

The Linguistics Enterprise

From knowledge of language
to knowledge in linguistics

Edited by

Martin Everaert

Tom Lentz

Hannah De Mulder

Øystein Nilsen

Arjen Zondervan

John Benjamins Publishing Company

The Linguistics Enterprise

Linguistik Aktuell/Linguistics Today (LA)

Linguistik Aktuell/Linguistics Today (LA) provides a platform for original monograph studies into synchronic and diachronic linguistics. Studies in LA confront empirical and theoretical problems as these are currently discussed in syntax, semantics, morphology, phonology, and systematic pragmatics with the aim to establish robust empirical generalizations within a universalistic perspective.

General Editors

Werner Abraham
University of Vienna /
Rijksuniversiteit Groningen

Elly van Gelderen
Arizona State University

Advisory Editorial Board

Josef Bayer
University of Konstanz

Cedric Boeckx
ICREA/Universitat Autònoma de Barcelona

Guglielmo Cinque
University of Venice

Liliane Haegeman
University of Ghent

Hubert Haider
University of Salzburg

Terje Lohndal
University of Maryland

Christer Platzack
University of Lund

Ian Roberts
Cambridge University

Lisa deMena Travis
McGill University

Sten Vikner
University of Aarhus

C. Jan-Wouter Zwart
University of Groningen

Volume 150

The Linguistics Enterprise.

From knowledge of language to knowledge in linguistics

Edited by Martin Everaert, Tom Lentz, Hannah De Mulder, Øystein Nilsen
and Arjen Zondervan

The Linguistics Enterprise

From knowledge of language
to knowledge in linguistics

Edited by

Martin Everaert¹

Tom Lentz¹

Hannah De Mulder¹

Øystein Nilsen²

Arjen Zondervan¹

¹Utrecht University / ²Queen Mary, University of London

John Benjamins Publishing Company

Amsterdam / Philadelphia



The paper used in this publication meets the minimum requirements of American National Standard for Information Sciences – Permanence of Paper for Printed Library Materials, ANSI Z39.48-1984.

Library of Congress Cataloging-in-Publication Data

The linguistics enterprise : from knowledge of language to knowledge in linguistics / edited
by Martin Everaert ... [et al.].

p. cm. (Linguistik Aktuell/Linguistics Today, ISSN 0166-0829 ; v. 150)

Includes bibliographical references and index.

1. Cognitive grammar. 2. Semantics. 3. Grammar, Comparative and general--Syntax.

I. Everaert, Martin.

P165.L54 2010

415--dc22

2009043585

ISBN 978 90 272 5533 4 (Hb ; alk. paper)

ISBN 978 90 272 8866 0 (Eb)

© 2010 – John Benjamins B.V.

No part of this book may be reproduced in any form, by print, photoprint, microfilm, or any other means, without written permission from the publisher.

John Benjamins Publishing Co. · P.O. Box 36224 · 1020 ME Amsterdam · The Netherlands
John Benjamins North America · P.O. Box 27519 · Philadelphia PA 19118-0519 · USA

Table of contents

List of contributors	VII
The linguistics enterprise: From knowledge of language to knowledge in linguistics	1
<i>Martin Everaert, Tom Lentz, Hannah De Mulder, Øystein Nilsen & Arjen Zondervan</i>	
Scope ambiguities through the mirror	11
<i>Raffaella Bernardi</i>	
Phonetic and phonological approaches to early word recognition: Empirical findings, methodological issues, and theoretical implications	55
<i>Paola Escudero & Titia Benders</i>	
Restructuring head and argument in West-Germanic	79
<i>Arnold E. Evers</i>	
Scope assignment in child language: On the role of the Question Under Discussion	99
<i>Andrea Gualmini & Sharon Unsworth</i>	
The learnability of A-bar chains	115
<i>Jacqueline van Kampen</i>	
Looking at anaphora: The psychological reality of the Primitives of Binding model	141
<i>Arnout W. Koornneef</i>	
Incremental discourse processing: How coherence relations influence the resolution of pronouns	167
<i>Willem M. Mak & Ted J.M. Sanders</i>	
Theoretical validity and psychological reality of the grammatical code	183
<i>Ad Neeleman & Hans van de Koot</i>	
Monitoring for speech errors has different functions in inner and overt speech	213
<i>Sieb Nooteboom</i>	
What's in a quantifier?	235
<i>Rick Nouwen</i>	

Minimal versus not so minimal pronouns: Feature transmission, feature deletion, and the role of economy in the language system <i>Eric Reuland</i>	257
Against partitioned readings of reciprocals <i>Sivan Sabato & Yoad Winter</i>	283
The representation and processing of fixed and compositional expressions <i>Joost Schilperoord & Rein Cozijn</i>	291
Clitic doubling in Spanish: Agreement of the third kind <i>Jan Schroten</i>	315
Metalinguistic processing and acquisition within the MOGUL framework <i>Michael Sharwood Smith</i>	327
Catching heffalumps: Onset fricatives in child speech <i>Wim Zonneveld</i>	345
Index	377

List of contributors

Raffaella Bernardi

Faculty of Computer Science
Free University of Bozen-Bolzano
Via della Mostra 4,
39100 Bolzano, Italy.
bernardi@inf.unibz.it

Titia Benders

Institute of Phonetic Sciences
University of Amsterdam
Spuistraat 210, 1012 VT Amsterdam
The Netherlands
titia.benders@uva.nl

Paola Escudero

Institute of Phonetic Sciences
University of Amsterdam
Spuistraat 210, 1012 VT Amsterdam
The Netherlands
escudero@uva.nl

Rein Cozijn

Department of Communication and
Information Sciences
Tilburg University
PO Box 90153, 5000 LE Tilburg
The Netherlands
r.cozijn@uvt.nl

Martin Everaert

Utrecht institute of Linguistics OTS
Utrecht University
Janskerkhof 13, 3512 BL Utrecht
The Netherlands
m.b.h.everaert@uu.nl

Arnold E. Evers

Utrecht institute of Linguistics OTS
Utrecht University
Janskerkhof 13, 3512 BL Utrecht
The Netherlands

Andrea Gualmini

Utrecht institute of Linguistics OTS
Utrecht University
Janskerkhof 13, 3512 BL Utrecht
The Netherlands
a.gualmini@uu.nl

Jacqueline van Kampen

Utrecht institute of Linguistics OTS
Utrecht University
Janskerkhof 13, 3512 BL Utrecht
The Netherlands
j.vankampen@uu.nl

Arnout W. Koornneef

Utrecht institute of Linguistics OTS
Utrecht University
Janskerkhof 13, 3512 BL Utrecht
The Netherlands
a.w.koornneef@uu.nl

Hans van de Koot

Department of Phonetics and Linguistics
University College London
Chandler House
2 Wakefield Street, London WC1N 1PF
United Kingdom
h.v.d.koot@ucl.ac.uk

Tom Lentz

Utrecht institute of Linguistics OTS
Utrecht University
Janskerkhof 13, 3512 BL Utrecht
The Netherlands
t.o.lentz@uu.nl

Willem M. Mak

Department of Dutch
Utrecht University
Trans 10, 3512 JK Utrecht
The Netherlands
w.m.mak@uu.nl

Hannah De Mulder

Utrecht institute of Linguistics OTS
Utrecht University
Janskerkhof 13, 3512 BL Utrecht
The Netherlands
h.n.m.demulder@uu.nl

Ad Neeleman

Department of Phonetics and Linguistics
University College London
Chandler House
2 Wakefield Street, London WC1N 1 PF
United Kingdom
a.neeleman@ucl.ac.uk

Øystein Nilsen

Department of Linguistics
Queen Mary, University of London
Mile End Road, London E1 4NS
United Kingdom
o.nilsen@qmul.ac.uk

Sieb Nooteboom

Utrecht institute of Linguistics OTS
Utrecht University
Janskerkhof 13, 3512 BL Utrecht
The Netherlands
s.g.nooteboom@uu.nl

Rick Nouwen

Utrecht institute of Linguistics OTS/
Department of Dutch
Utrecht University
Janskerkhof 13, 3512 BL Utrecht
The Netherlands
r.nouwen@uu.nl

Eric Reuland

Utrecht institute of Linguistics OTS
Utrecht University
Janskerkhof 13, 3512 BL Utrecht
The Netherlands
e.j.reuland@uu.nl

Sivan Sabato

School of Computer Science and
Engineering,
The Hebrew University-Givat-Ram
Edmond Safra Campus
Jerusalem, 91904
Israel
sivan_sabato@cs.huji.ac.il

Ted J. M. Sanders

Department of Dutch
Utrecht University
Trans 10, 3512 JK Utrecht
The Netherlands
t.j.m.sanders@uu.nl

Joost Schilperoord

Department of Communication and
Information Sciences
Tilburg University
PO Box 90153, 5000 LE Tilburg
The Netherlands
j.schilperoord@uvt.nl

Jan Schroten

Utrecht institute of Linguistics OTS
Utrecht University
Janskerkhof 13, 3512 BL Utrecht
The Netherlands

Michael Sharwood Smith

Department of Language and Intercultural
Studies
Heriot-Watt University
EH14 4AS Edinburgh
United Kingdom
m.a.sharwood_smith@hw.ac.uk

Sharon Unsworth

Utrecht institute of Linguistics OTS/Meertens
institute
Utrecht University
Janskerkhof 13, 3512 BL Utrecht
The Netherlands
s.unsworth@uu.nl

Yoad Winter

Department of Modern Languages
Utrecht University
Trans 10, 3512 JK Utrecht
The Netherlands
y.winter@uu.nl

Wim Zonneveld

Department of Modern Languages
Utrecht University
Trans 10, 3512 JK Utrecht
The Netherlands
w.zonneveld@uu.nl

Arjen Zondervan

Utrecht institute of Linguistics OTS
Utrecht University
Janskerkhof 13, 3512 BL Utrecht
The Netherlands
a.j.zondervan@uu.nl

The linguistics enterprise

From knowledge of language to knowledge in linguistics

Martin Everaert, Tom Lentz, Hannah De Mulder,
Øystein Nilsen & Arjen Zondervan

1. Introduction

The study of language and speech explores the interfaces between the rules and laws underlying human language, the systems of grammar, and the manner in which these interact with internal systems and processes (interpretation, speech perception and production) and with the outside world (acquisition, use, change and role in society). The study of language thus defined, centres on three dimensions, summarised in the following (partly subsequent) questions:

- i. What constitutes knowledge of language? The study of the architecture of the language system, i.e. language taken as a cognitive system, the computational/ logical modelling of this system, and the study of the interaction between its various components.
- ii. How is language acquired? The study of the cognitive processes underlying primary and secondary language acquisition.
- iii. How is language put to use, and used? The linguistic processes underlying language production, perception and interpretation under varying conditions, such as modality, social environment and speech context.

In this introduction, we will apply these questions to linguistics itself by abstracting from *knowledge of language* to *knowledge in linguistics*. In Section 2, we will address the question what knowledge in linguistics is by presenting an overview of the goals in current linguistic research. In Section 3, we will discuss how knowledge of language is acquired in linguistics, by discussing the range of techniques used by linguists. Finally, in Section 4, we will consider how the knowledge of the language system that is acquired by linguists is put to use to gain more in-depth knowledge of human cognition.

2. What is knowledge in linguistics?

2.1 The linguistics enterprise

Linguistics investigates the systems underlying language, speech and language use. Linguists seek to develop an understanding of the rules and laws that govern the structure and use of particular languages. They search for the general laws and principles governing all natural languages, i.e. the nature of the computational system of human language in its many guises and the way it fulfils our communicative needs. In other words, linguistics aims at a deeper understanding of the constraints and forces that shape human languages.

Linguists study all aspects of language: from the generation of speech sounds and their acoustical properties, to the role of language in social cohesion, and to how language gets processed by the brain and interpreted. To do so, we need a thorough understanding of the architecture of the language system and in-depth knowledge of its main components.

There are many approaches to the study of language. Some linguists study language as an independent system (architecture), describing and showing the interaction between the various parts of that system (phonetics and phonology, morphology and the lexicon, syntax, semantics and pragmatics). Other researchers are concerned with the processes of listening, reading, speaking and writing, and the acquisition of language (psycholinguistics, neurolinguistics, language acquisition). One might also study language as a cultural phenomenon, as a tool for social interaction, or one might take a historical perspective and study the familial relations among languages or how languages have changed over time.

In the last decades, one particular approach to the study of language has been prominent. According to this approach, language is an interesting phenomenon as it represents a structured and accessible product of the human mind, a window into the cognitive abilities of man: language studied as a means to understand the nature of the mind that produces it. Linguistics, thus, became part of cognitive science, the study of the higher-level mental functions, including memory, reasoning, visual perception/recognition, etc. It is this approach that is the starting point for theoretical frameworks such as Generative Grammar (in all its manifestations), Cognitive Linguistics, and many functional approaches.

Our ability to use language to communicate is a complex skill. It requires the ability to structure knowledge, to encode it in linguistic signals, and to interact with others using such signals. The acquisition of a language is, likewise, a highly complex process. Virtually all children learn their native languages quickly, with most of the crucial elements firmly in place by the age of six. Adults exhibit quite a different path when they acquire a second language, and do not necessarily come close to the

level of native speakers, even after a much longer time of learning and/or exposure. Crucial to present day linguistics is the question how children are able to acquire the intricate principles and properties of the language of their environment in such a short amount of time. Most researchers would agree that children acquire language through an interplay of innate mechanisms (for some a language module) and environmental factors. This poses the interesting question how nature and nurture interact in language learning. Additionally, the nature of innateness has to be studied. Can the ability to learn a language be explained by assuming innate cognitive structures that guide learning in a certain direction, or are specific elements of language itself hard-wired into the brain?

2.2 Subdisciplines in linguistics

The rapid growth of the field has led to a certain degree of fragmentation of linguistics into subdisciplines, between which cross-fertilisation has become increasingly rare. Over the years, the discipline has become more and more a conglomerate of largely independent research areas, such as sociolinguistics, syntax, language acquisition, discourse studies, phonology, computational linguistics, pragmatics, phonetics, etc. This is reflected in the existence of journals with names like “Journal of X” or “X” where X stands for the aforementioned list of subdisciplines. This is an inevitable development in every field, but at the same time it is unfortunate.

It is important to keep the discussion between the subdisciplines alive and keep in mind what the underlying, more fundamental questions about language are that we want to address as linguists. Focusing on different aspects of the questions that drive all linguists does not need to lead to mutual exclusion, but rather to complementary results. This volume tries to reflect that unified approach to the study of language. In this volume contributions from different subdisciplines are brought together: computational linguistics (Bernardi), discourse (Mak and Sanders; Nouwen; Schilperoord and Cozijn), language acquisition (first language acquisition: Escudero and Benders; Gualmini and Unsworth; Van Kampen; Zonneveld; second language acquisition: Sharwood Smith), phonetics (Escudero and Benders; Nooteboom), phonology (Escudero and Benders; Zonneveld), psycholinguistics (Koornneef; Neeleman and van de Koot), semantics (Nouwen; Sabato and Winter), syntax (Evers; Neeleman and van de Koot; Reuland; Schroten). At an empirical level, the papers focus on: anaphoric interpretation (Koornneef; Mak and Sanders; Reuland), clitics (Schroten), verb clustering (Evers), connectives (Mak and Sanders), idioms (Schilperoord and Cozijn), quantifier (scope) (Bernardi; Nouwen; Gualmini and Unsworth), speech errors (Nooteboom), *wh*-elements (Van Kampen), word recognition (Escudero and Benders) and onset fricative avoidance (Zonneveld).

3. How do we acquire knowledge in linguistics? Linguistic methodology

3.1 Experimental methods

The methods used in the study of language are diverse and can be roughly divided into three groups: behavioural, computational, and neurophysiological, such as electro-physiological measures (EEG), event-related potentials (ERP), magnetoencephalography (MEG) and magnetic resonance imaging (MRI)). It depends on the subdiscipline one is working in which methods are preferred. Do we, for instance, need psychological evidence to show which grammar's rules are 'psychologically real'? By that we mean, will only evidence acquired by perceptual tests, such as self-paced reading, cross-modal priming, visual-world paradigm, or tracking of eye movements count? No, any evidence will do, whether acquired by experimentation, or the use of corpora, or by assessing linguistic structures with our own judgements. As long as one can show that the data bear relevance to one's hypothesis, no methodology has any a priori preferential status. It depends on the questions one has which research methodology is appropriate. In this volume the type of research reported on is based on eye-tracking experiments (Koornneef; Schilperoord and Cozijn), reaction-time experiments (Escudero and Benders), elicited speech errors (Nooteboom), the use of corpora (Schilperoord and Cozijn; Mak and Sanders), longitudinal data (Van Kampen; Zonneveld), truth value judgment tests (Gualmini and Unsworth), intuitive judgements (Bernardi; Evers; Reuland;¹ Schroten; Sabato and Winter) and questionnaires (Nouwen).

The use of intuitive judgments has long been the hall-mark of theoretical research in syntax, semantics, and to a lesser degree in morphology and phonology.² We think that that doesn't hold true any more for much grammatical work now. If we take the work that is done in the institutions we are presently working at – Utrecht University and Queen Mary, London – we see syntacticians, semanticists, phonologists working together with psycholinguists, acquisitionists and sociolinguists, using experimental work to support their theoretical claims, making use of intuitive judgements, sometimes supporting them with magnitude estimation experiments, but similarly making use of corpora and questionnaires.

3.2 Theoretical and computational modelling

We use models to help us make sense of reality, or attend only to those features of it that are of interest to us in our enquiry. To model language is to provide an abstract

1. But do note that Reuland crucially bases his work on psycholinguistic and neurolinguistic work of others.

2. In a recent paper Wasow and Arnold (2005) claim that that situation still holds true for what they call 'generative grammarians'.

representation of its central characteristics, a grammar, so that it becomes easier to see how it is structured and used. Linguists cannot make much progress without using models, but they will have strongly diverging ideas about which models are appropriate for their aims.

A generative linguistic theory attempts to describe the tacit knowledge of a native speaker through a system of rules that, in an explicit and well-defined way, specifies all and only the well-formed sentences of a language. The study of the acquisition of language, for instance, must be solidly grounded in the study of the internal structure of the grammar. We need a thorough understanding of the architecture of the language subsystems and how the different components of grammar and the systems of use interact. The papers of Neeleman and van de Koot and Sharwood Smith sketch such a more comprehensive architecture.

Such understanding cannot be achieved without in-depth study of its main components: phonetics and phonology, morphology and the lexicon, syntax, semantics and pragmatics. The papers of Bernardi, Evers, Nouwen, Reuland, Schroten and Sabato and Winter take this route, discussing a certain phenomenon in detail, allowing us to build on that in language acquisition research, psycholinguistics and discourse studies. To some extent, computational modelling has also been used to simulate how language may be processed by the brain to produce behaviours that are observed. Computer models have been used to address several questions in parsing as well as how speech information is used for higher-level processes, such as word recognition.

4. How do we put knowledge in linguistics to use? Capturing cognition

Traditionally, linguistics was firmly rooted in the research traditions of the humanities. Studying language and speech was not only possible through curricula offered by departments of (comparative) linguistics/phonetics, but also in programmes offered by departments of modern language and culture, such as the department of English, the department of Spanish, the department of Dutch etc. In the last decades of the 20th century, the study of language moved away from its philological roots, positioning itself more and more as a cognitive science. Nowadays, the study of language and speech is predominantly seen as part of the Cognitive Sciences, constituting a natural federation of scientific disciplines, including psychology, neuroscience, artificial intelligence, and information science, with a shared research agenda.

Underlying the cognitive sciences is the conception of human beings – individually as well as collectively – as information processing systems. This deep analogy makes it possible for researchers to study natural and artificial intelligence, including the interactions between the two, from a coherent theoretical perspective. This conception has fundamentally changed the outlook of linguistics by establishing close connections both with the formal sciences, leading to computational linguistics, and

with psychology, resulting in the joint venture of psycholinguistics. The recent efforts to link linguistics to biology and cognitive neuroscience hold considerable promise for a further dimension in the understanding of the cognitive faculties involved in language. This does not entail a reduction of linguistics to other disciplines. It depends on one's position on the question whether mental entities like consciousness are reducible to the algorithm or system that produces them. Language philosophers like Searle posited this problem (the Chinese Room argument (Searle 1980)). This does not mean, however, that there is no viable enterprise in connecting linguistic theory with psychology, mathematics and computer science, biology, or neuroscience. Many studies, including a number of contributions to this volume, are examples of that fact.

We will address the question of how knowledge of the language system is obtained in linguistics relates to human cognition by discussing three important topics in the current debate on this issue: psychological reality, modularity and the competence/performance distinction.

4.1 Psychological reality

The introduction of generative grammar (re)oriented linguistics towards the 'psychological reality' of grammars, as an object of study instead of a mere part of the black box that produces (linguistic) behaviour. Psychological reality is usually meant to indicate an underlying claim about what such a model is attempting to explain. Psychological reality is a claim about the validity of a theory, where validity is conformity to or correspondence with external reality. Thus, the theory of Universal Grammar is psychologically real and the principles it incorporates are valid to the extent that they conform to external reality. Following Smith (2004), we assume that, to a large extent, such a conception is uncontroversial. Many would agree that our linguistic behaviour can be described in terms of rules and principles. Paraphrasing Smith (op. cit p.97), it seems unexceptionable to state that we have something like a 'rule' that states that determiners precede nouns and verbs precede particles in English – *the man/*man the; buy a car/*a car buy* – but it is often conceived as quite a different matter whether there is a rule that moves a *wh*-element to sentential-first position, leaving behind a trace (*Who_i did you see e_i*). Postulating knowledge 'in our head' of categories of the type determiner, noun, verb, particle and the concept of left-right ordering seems uncontroversial, but claiming the psychological reality of the concept of movement of categories, leaving behind traces, is often considered controversial. It is simply an empirical hypothesis that can be proven wrong, or right, for that matter (cf. Friedmann et al. 2008).

The topic of psychological reality is addressed in a number of papers in this volume. Koornneef looked into the psychological reality of the Primitives of Binding model of Reuland (2001) by doing eye-tracking experiments to show that the economy hierarchy at the basis of Reuland's model is reflected in real time anaphora

comprehension. Gualmini and Unsworth do the same for two other theoretical constructs, surface scope and Question Under Discussion. Nouwen tries to reconcile what linguists know about quantifying expressions with what cognitive psychologists have observed. He tries to bring together the psychological models of quantifier meaning and the models provided by formal semantic theory. Van Kampen observes that for the study of language acquisition it is evident that grammatical distinctions should reflect psychological reality, but she also observes that it is not immediately clear whether our research has advanced enough to make such claims tangible. Sharwood Smith argues that his model of (second) language acquisition not only explains linguistic knowledge in the individual and the logical problem of language acquisition, but also accounts for how language is processed on-line.

4.2 Modularity

A long standing question in cognitive science is how we should think about the structure of the mind. To what extent does the mind/brain consist of separate faculties that are each specific to particular types of processes? We can consider different kinds of answers to this question: all of the mind could be recruited in all of the things we do or specific areas of the mind could be dedicated to particular cognitive domains. The idea that particular cognitive tasks are dealt with by specific areas in the mind is termed the modularity theory of mind, discussed at length in Fodor (1983). However, Fodor was not the first to consider the mind/brain as constituting of different parts: in the 18th century, phrenologists already proposed that the different brain areas were engaged in different cognitive functions.

But what does it mean for a particular part of the brain to be engaged in a specific cognitive function? According to Fodor, a particular cognitive function corresponds to a module in the brain if a number of characteristics have been met. Among other things, the module for a particular cognitive function has to be domain specific, mandatory, fast and informationally encapsulated. An example will make these characteristics more concrete: take the cognitive function vision. In order to determine whether vision is a module in the Fodorian sense, the vision module should respond only to visual stimuli, not to any other types of stimuli. If so, it can be considered domain specific. In order to be considered mandatory and fast, the vision module has to “see” something rapidly whenever it gets the appropriate visual input, whether or not it makes sense in the given context. The vision module should immediately register a tiger when it gets the appropriate visual input for tigers, even though the tiger is seen in a highly unlikely context (a shopping mall in the Netherlands, for instance). This last point also relates to the notion of informational encapsulation: a module cannot refer to the output or workings of other areas of the brain not part of the module. The module can thus only work with its specified input, which, in the case of the example, is visual stimuli.

Aside from vision, many other cognitive functions are also considered to be modules in this sense. The relevant question for our present purposes is: is language one of them? In other words: does the language module only respond to language? Is it mandatory that we parse language when we hear it and do we do so quickly? And do we recruit knowledge from outside the direct linguistic input when we use language? Especially this last question has been the topic of much debate in linguistic inquiry over the last few decades. Is language informationally encapsulated? If it is, we must assume that world knowledge, hearer's expectations, speaker's intentions etc. do not play a role in the processing and production of language. An alternative would be to assume that only particular aspects of language (grammatical processing is a prime candidate) are informationally encapsulated.

Issues of the kind described above are addressed in many of the contributions to this volume. In his study, Koornneef, for example, assumes that syntax and discourse are separate modules of language. In his view, syntactic operations in the domain of anaphora resolution are fast and mandatory and they are not influenced by available discourse information. An opposite view is presented in the contribution by Mak and Sanders. They claim that contextual and discourse information given by the use of particular causal connectives does influence syntactic processing. Aside from these two contributions that address this issue quite explicitly, the role of extra-linguistic knowledge in language processing is also a running theme in a number of the other contributions. Sabato and Winter, for example, consider the role of contextual factors on the partitioned reading of reciprocals. In the domain of second language acquisition, Sharwood Smith's contribution concerns the influence of metalinguistic knowledge on the acquisition of language. His central question is whether language learners are able to recruit metalinguistic knowledge to aid them in the language acquisition process. Regarding first language acquisition, Gualmini and Unsworth consider the role of the discourse in children's ability to resolve ambiguity in scope assignment.

Contributions from different fields of linguistics (first and second language acquisition, syntactic processing, discourse analysis) are thus all concerned with the modular status of language in general and its informational encapsulation status in particular.

4.3 The competence/performance distinction

In any field of study the question arises whether we should investigate the object of study by taking external factors and imperfect behaviour into account, or whether we should aim at uncovering idealised underlying laws, abstracting away from irrelevant interfering factors in the behaviour of the object of study. As Niaz (1999) points out, an important breaking point in this discussion in the natural sciences was when Galileo presented his law of free fall, which idealised away from interfering factors such as air resistance, something many of his precursors in the Middle Ages

had considered unacceptable. A key issue in the debate was that the idealised law can never be observed in reality, but has to be extrapolated from observations that approach the ideal situation.

This 'Galilean idealisation' made its way into the psychological sciences when Piaget made a distinction between *epistemic subject* and *psychological subject*, the first one being the idealised version of the latter (see Niaz 1999). The epistemic subject is the 'ideal knower', the underlying competence of subjects, ignoring irrelevant limitations that affect performance in experiments. The rationale is that just like in the natural sciences, the underlying competence can be approximated by varying the irrelevant interfering factors and extrapolating from those findings.

This distinction between underlying *competence* and observable *performance* was applied to linguistics most famously by Chomsky (1965). He posited that the subject matter of linguistics should be the mental system of rules that makes up the speaker's knowledge of language, not the behaviour that results from application of these rules in production and comprehension. However, just like in Galileo's case, the competence is only observable by observing performance, so the same method of extrapolation to the ideal has to be applied.³

Conceptually, there are two alternatives in linguistics to this method of approximation of the underlying ideal. The first one is to break with the idea that the subject matter of linguistics should be only the idealised mental system. Many linguists are interested in the way the language system works when factors that Chomsky (1965) would consider interfering, are taken to be part of the language system itself. For instance, much research is done on the interplay of memory and mechanisms of language production and comprehension (e.g. parsing), while one could decide to consider memory limitations to be a performance factor that lies outside the language system. The most radical version of this view of linguistics would be to investigate only performance itself, without any idealisation to speakers' competence. However, the question then arises whether any meaningful generalisations can be made without reference to an idealised competence, or whether, as Chomsky puts it: 'investigation of performance will proceed only so far as understanding of underlying competence permits' (1965: 10).

The second alternative in linguistics to idealisation from performance to competence is to focus on competence only. This can be done by taking the data that is accounted for by an existing theory of competence and designing a new theory of competence that accounts for the same data (it has the same *descriptive* adequacy),

3. Chomsky points out that even when intuitive judgments by native speakers are used, performance factors may interfere and that it is the task of the linguist to extract these intuitions in such a way that extrapolation to the idealised competence is possible.

but in a manner that is preferable to the existing theory (it has higher *explanatory adequacy*). This preference can be defined on many dimensions: economy principles, neurological considerations, or (Chomsky's favourite): preference from the perspective of language acquisition. The question this approach faces is when to stop looking for more preferable theories that cover the same data.

This topic is explicitly addressed by Neeleman and Van de Koot's paper. Instead of taking the competence-performance distinction as reflecting that grammar is a knowledge base consulted by the performance systems, they argue for an alternative interpretation. In their approach the grammar and the performance systems are theories of the same object, but at different levels of description. The articles of Mak and Sanders, Nooteboom, Schilperoord and Cozijn, and Sharwood Smith are focussed on actual language use, and it is not immediately clear whether a competence-performance distinction is meaningful in their area of research.

In whatever way the articles in this volume assess the relation of language to human cognition, they all present a valuable addition to answering the three questions we raised above about knowledge in linguistics: What is knowledge in linguistics, how is it acquired and how is it put to use? The reader can of course use this book to look up data on the specific phenomena addressed in the articles, but we hope it will also present insights into the palette of present-day linguistics. In this way, we hope to break open the division of linguistics into subfields and to make cross-fertilisation possible.

References

- Chomsky, N. 1965. *Aspects of the Theory of Syntax*. Cambridge MA: The MIT Press.
- Fodor, J.A. 1983. *The Modularity of Mind*. Cambridge MA: The MIT Press.
- Friedmann N., Taranto G., Shapiro L.P. & Swinney, D. 2008. The leaf fell (the leaf): The online processing of unaccusatives. *Linguistic Inquiry* 39: 355–377.
- Niaz, M. 1999. The role of idealization in science and its implications for science education. *Journal of Science Education and Technology* 8(2): 145–150.
- Reuland, E. 2001. Primitives of binding. *Linguistic Inquiry* 32: 439–492.
- Searle, J. 1980. Minds, brains, and programs. *Behavioral and Brain Sciences* 3(3): 417–457.
- Smith, N. 2004. *Chomsky: Ideas and Ideals*. Cambridge: CUP.
- Wasow, T. & Arnold, J. 2005. Intuitions in linguistic argumentation. *Lingua* 115: 148–1496.

Scope ambiguities through the mirror

Raffaella Bernardi*

Free University of Bozen-Bolzano

In this paper we look at the interpretation of Quantifier Phrases from the perspective of Symmetric Categorical Grammar. We show how the apparent mismatch between the syntactic and semantic behaviour of these expressions can be resolved in a typological system equipped with two Merge relations: one for syntactic units, and one for the evaluation contexts of the semantic values associated with these syntactic units.

1. The logic behind natural languages

“The Mathematics of Sentence Structures”, the title of (Lambek 1958), clearly summarizes the ambitious challenge Jim Lambek has launched to researchers interested on formal approaches to natural language. It could also be read as asking “which is the logic that captures natural language composition?”, which can be rephrased as “which are the models of natural language composition?”, “which are its formal properties?”.

Since (Montague 1974), the λ -calculus was used by linguists to build the meaning representation of sentences compositionally. The high expectations on the line of research pioneered by Lambek were even increased by Johan van Benthem’s results (Van Benthem 1987) on the correspondence between the identified logic (known as Lambek calculus) and the λ -calculus. However, though all the questions above, and even more, were answered pretty soon (see (Lambek 1961; Moortgat 1988) among others), the answers found hold only for a fragment of natural language. What remains outside of the defined logic (and hence model and proof system) seemed to be those structures containing non-local dependencies and scope operators. The proof of the

*I would like to express my gratitude to Chris Barker, Marco Baroni, Anna Szabolcsi and Roberto Zamparelli for their valuable feedback on earlier versions of the present paper. My thanks go to Chris Barker, Philippe de Groote, Glyn Morrill, Richard Moot, Ian Pratt, Ken Shan, and Mark Steedman for stimulating discussions on the QP puzzle, and on the approach to natural language analysis undertaken in this paper. Finally, the work reported in this paper is part of the adventure started with Rajeev Goré and Michael Moortgat back at ESSLLI 2004 and brought ahead with Michael Moortgat through all these years. The work reported has been carried out jointly with Michael Moortgat.

lack of expressivity of the Lambek calculus was given by the results on the correspondence of the proposed logics with Context Free Language, since it is known that natural language is outside this class of formal languages. Through the years, many proposals have been made to overcome these well defined limitations, but these proposals departed in a way or another from the original line of research Lambek had launched, e.g. they strive for wide coverage of the parser rather than for logical completeness of the system (Steedman 2000).

In this paper, we go back to 1961 and start from there again, but with the new knowledge put at our disposal by research on algebra, logic, programming languages and linguistics. We show that the time is ready to recover the original enthusiasm on the possibility to solve Lambek's original challenge. In particular, we focus on scope operators and the mismatch that they might create between syntactic and semantic structures. The original questions regarding the model, the proof systems, and the meaning representation of the structures recognized by the logic we will present are formally answered in (Kurtonina & Moortgat 2007; Bernardi & Moortgat 2007; Moortgat 2007; Bernardi and Moortgat to appear, Moortgat 2009); here we give a more gentle introduction to the system we propose and focus on its contribution to the linguistic research on scope in natural language.

Lambek brought to the light the mathematical apparatus behind Categorical Grammar (CG) showing that CG was based on half of the residuation principle and illustrated the advantages that one would gain by considering the whole mathematical principle given below.

$$(1) \quad (\text{RES}) \quad A \subseteq B/C \text{ iff } A \otimes C \subseteq B \text{ iff } C \subseteq A \backslash B$$

In more familiar term, the Lambek calculus, or calculus of residuation, consists of function application and abstraction rules, which are used to prove the grammaticality of a given sentence. For instance, given that *Alex* belongs to the category of determiner phrases (dp) and *left* to the one of intransitive verbs ($dp \backslash s$) (i.e. a function missing a dp on the left to yield a sentence s), the expression *Alex left* is proved to belong to the category s simply by function application:

$$(2) \quad \begin{array}{lll} \text{If} & Alex & \in dp \\ & left & \in dp \backslash s \\ \text{Then} & Alex \text{ left} & \in s \end{array}$$

which is the proof of the following theorem:

$$(3) \quad (a) \quad dp \otimes (dp \backslash s) \vdash s$$

While function application was already used in CG, Lambek showed that the residuation calculus includes also abstraction which is needed, for instance, to prove the statement below, i.e. the type lifting rule used by linguists.

$$(4) \quad (b) \quad dp \vdash s / (dp \backslash s)$$

We propose to extend the expressivity of the logical grammar by moving to a symmetric calculus, i.e. a calculus in which each rule has a dual. The concept of *dual rules* and *dual theorems* was introduced by Schröder (Schröder 1980) who noticed that the rules of propositional logic could be formulated in pairs corresponding to De Morgan dual operators, e.g. the rules about \wedge are dual to rules about \vee and so are the theorems below involving them:

$$(5) \quad A \wedge (B \vee C) \vdash (A \wedge B) \vee (A \wedge C) \text{ dually } (C \vee A) \wedge (B \vee A) \vdash (C \wedge B) \vee A$$

The logic language of the Lambek calculus consists of residuated operators $((A \backslash B), (A \otimes B), (B / A))$. We extend it with their duals, viz. $((B \oslash A), (B \oplus A), (A \odot B))$ and claim that besides the residuation principle (RES) the mathematics of sentence structures includes its dual principle (DRES),

$$(6) \quad (\text{DRES}) \quad B \oslash A \subseteq C \text{ iff } B \subseteq C \oplus A \text{ iff } C \odot B \subseteq A$$

together with the properties governing the interaction between residuated and dual residuated operators. These interaction properties were first studied by V.N. Grishin (Grishin 1983). We call the new system Lambek Grishin (LG) (See Section 3).

In the Lambek calculus the theorems $dp \otimes dp \vdash s$ and $dp \vdash s / (dp \backslash s)$ are proved by function application and abstraction, in the extended system also their duals hold:

$$(7) \quad (a') \quad s \vdash (s \odot dp) \oplus dp \quad (b') \quad (s \odot dp) \odot s \vdash dp$$

are proved by co-application and co-abstraction, i.e. the dual rules of application and abstraction.

In this new framework, we can distinguish categories standing for linguistic expressions and categories standing for the contexts of those expressions: contexts of category A are structures with an A -hole. The \otimes -relation merges linguistic expressions and the \oplus -relation merges contexts where these expressions can be plugged in. For instance, $dp \backslash s$ stands for *left* and *knows mary*, whereas $s \odot dp$ stands for the contexts of *left* and *knows mary*, viz. the tree with a hole for that expression; for atomic formulas, we should distinguish between e.g. dp , standing for *mary*, and dp^c standing for the context of a dp . We can avoid marking this difference at the level of atomic formulas since the merging relation used disambiguates them.

We briefly discuss the Quantifying In mechanism proposed in (Montague 1974) and its variation, Quantifier Raising, proposed by linguists to describe the behavior of quantifiers (May 1977; Reinhart 1997) (Section 4). The logic obtained is in the Curry Howard Correspondence with the $\bar{\lambda}\mu\tilde{\mu}$ -calculus. The meaning representation of the parsed structures are built thanks to this correspondence. For reason of space, in this paper we won't discuss this part; the reader interested in the issue is referred to (Bernardi & Moortgat 2007, Bernardi and Moortgat to appear); the only aspect to highlight here is that this extended lambda calculus has a way to represent and compose "context" too and to shift from one merging relation to the other.

In sum, the main novelty of our proposal is to introduce symmetry in categorial grammar by extending the Lambek calculus with the duals of its rules and with communication principles governing the interaction between the Lambek calculus and its dual. The new framework sheds new lights that we believe could be of interest for other frameworks too:

1. besides having categories standing for words, there is the need of having categories standing for the words' contexts;
2. hence, besides the merging relation, that builds a new syntactic unit from adjacent syntactic objects, natural language structures need also its dual relation, viz. the relation that merges contexts. (See (63) and (67) in Section 3.) Therefore, we can distinguish syntactic trees and context trees, the two types of trees can interact but their structure should not be modified once built. (See (70) in Section 3.)
3. non-local scope operators, like quantifiers, consist of a syntactic and a semantic component tied to each other. Thanks to the interaction between the two merging relations, the syntactic component stays in the syntactic structure or constituent-domain (c-domain) while the semantic component can jump out of it to reach the semantic-domain (s-domain), viz. the structure on which the QP has scope. (See (56) in Section 2 and Section 4.)
4. the c-domain is governed by the merging relation, whereas the s-domain is governed by the relation governing contexts. (See (82) and (95) in Section 4.)
5. a syntactic structure containing scope operators can be built in different ways: the order in which scope operators are activated determine the readings of the structure. (See (87), (91) in Section 4.)

2. Preliminary background

Logic is about entailment of a conclusion from some premises. We might want to know whether it holds that whenever it is true that *If Alex swims then Alex gets wet* ($p \rightarrow q$) and that *Alex swims* (p) then it is also true that *Alex gets wet* (q), which is formalized by the entailment below.

$$(8) \quad \{p \rightarrow q, p\} \vdash q$$

But as the objects of the reasoning vary, so it may vary the logic one needs. Our objects are natural language structures; hence we are interested in (a) tailoring the logic of natural language syntactic composition and (b) capturing its correlation with the semantic representation. We want to grasp the fragment of logic that suits our needs, and use a logic as a grammar. This approach is known as the “parsing as deduction” approach.

Since the number and order of occurrence of formulas (words) and their structure are important when dealing with natural language the premise of our entailment is not a set as in propositional logic but a structure. In propositional logic, both $\{p \rightarrow q, p\} \vdash q$ and $\{p \rightarrow (p \rightarrow q), p\} \vdash q$ hold, whereas we need a logic sensitive to the number and the order of the formulas.

As for (a), employing a logic as a grammar to analyse natural language expressions means to assign atomic or complex formulas to words and let their logical rules check that a given string of words is of a certain category and return its syntactic structure. Given the lexical entries $w_1 \dots w_n$, the system will take the corresponding syntactic categories $A_1 \dots A_n$ and prove which structures they can receive so as to derive a given category B , i.e. check whether a tree-structure of category B can be assigned to the given string of words.

$$(9) \quad \underbrace{A_1}_{w_1} \dots \underbrace{A_n}_{w_n} \vdash B$$

To be more precise, syntactic categories are sets of expressions: those expressions that belong to a given category, i.e. dp 's etc. are category labels: names for such sets. In short, we are dealing with e.g. Figure 1.

$$\underbrace{\{\text{john, mary}\}}_{np} \otimes \underbrace{\{\text{left, knows lori}\}}_{np \setminus s} \vdash \underbrace{\{\text{john left, mary left, john knows lori, mary knows lori}\}}_s$$

Figure 1. Example of syntactic categories as sets of expressions

We take the liberty to use the term “category” to refer to a category label. Similarly, the structure built out of the syntactic categories is a name for the sets of phrases that belong to that structure.

Since the introduction of the *Chomsky Hierarchy*, the attention of formal linguists has focused on understanding where natural languages fit in the hierarchy and hence, which formal grammars have the proper generative power to handle natural language structures. Similarly, when assuming a logical perspective the main question to address is which logic has the right expressivity to reason on natural language structures avoiding both under-generation and over-generation problems. In this perspective, the first question to ask is “which operators do we need to reason on natural language resources?”. Since Frege, the idea of interpreting words as functions and their composition as function application has been broadly accepted in Formal Semantics. Hence, implication (\rightarrow) is an essential connective to have in our logical language. Yet, it remains to be understood which properties it should have, which other operators we would need and how they relate to each other.

A well known approach to our second goal (b) is the assumption of a Logical Form (Chomsky 1976): the surface (syntactic) structure is input to a further set of rules which derive the Logical Form structure of the sentence. This latter structure is then interpreted by a semantic mechanism to obtain the meaning representation. Another well known solution is Montague's (Montague 1974) where syntax and semantics go in parallel: for each syntactic rule there is a corresponding semantic rule that builds the meaning representation of the phrase analyzed syntactically. The correspondence is postulated for each rule. This step-by-step procedure in some intermediate steps, e.g. in the quantification step, may build meaning representations also for structures which are not syntactic constituents. Our approach is similar to Montague's in treating syntax and semantics in parallel, but it differs from it into two ways. A first important difference with Montague's approach is that the correspondence between syntactic and semantic rules is defined at a general level: following (Van Benthem 1988), we exploit the correspondence introduced by Curry and Howard (Howard 1980) between logical rules and lambda calculus; the proof of the derivation $A_1, \dots, A_n \vdash B$ can be read as an instruction for the assembly of proof-term M with input parameters x_1, \dots, x_n . In the "parsing as deduction" approach this turns out to mean that *labelled derivations* provide also the schema to build the meaning representation of the structure under analysis:

$$(10) \quad x_1 : \underbrace{A_1}_{w_1} \dots x_n : \underbrace{A_n}_{w_n} \vdash M : B$$

It's important to emphasize that the correspondence holds between syntactic structures and semantic labels (proof terms). Recall the observation above that syntactic structures are labels for the set of expressions that belong to that structure. Once we instantiate the structure with a particular expression, we can also replace the semantic labels with the meaning representation of the words in the chosen expression. A second difference concerns the labelling of nodes: in LG only the syntactic constituents receive a semantic label.

Finally, there can be more than one way to prove that an entailment holds: different (normalized) derivations correspond to different readings of the same string of words. A given string can receive different structures and different proof terms, but also different proof terms for the same structure. In this paper, we are interested in the latter case (see Section 4). For reason of space, we won't discuss how a proof term is built by labelled derivations; however, since the linguistic phenomenon, we are interested in, lays on the interface of syntax and semantics of natural language, we introduce those concepts of Formal Semantics that are going to play a role in our analysis of quantifier phrases. The interested reader is referred to (Van Benthem 1987) for an introduction to the use of labelled derivations for building meaning representation of parsed

Table 1. Domain of interpretation

$[[\text{lori}]]$	=	lori;
$[[\text{alex}]]$	=	alex;
$[[\text{sara}]]$	=	sara;
$[[\text{pim}]]$	=	pim;
$[[\text{student}]]$	=	{alex, sara, lori};
$[[\text{teacher}]]$	=	{pim};
$[[\text{sing}]]$	=	{alex, sara, lori, pim};
$[[\text{trust}]]$	=	{⟨alex, alex⟩};
$[[\text{know}]]$	=	{⟨lori, alex⟩, ⟨lori, sara⟩ ⟨lori, lori⟩};
$[[\text{left}]]$	=	{lori};
$[[\text{every student}]]$	=	$\{Y \subseteq E \mid [[\text{student}]] \subseteq Y\}$.

strings, and to (Bernardi & Moortgat 2007, Bernardi and Moortgat to appear) for the discussion of the Curry-Howard correspondence for the logic presented in this paper with $\bar{\lambda}\mu\bar{\mu}$ -calculus and its translation into λ -terms.

2.1 Syntax-semantics interface in natural language

In order to illustrate how words are interpreted in Formal Semantics, let us consider the scenario below. We are in a school and the following people are present: *Lori*, *Alex*, *Sara* and *Pim*. Assume that the first three are students whereas *Pim* is a teacher. They all sing, *Alex* is the only one who trusts himself, *Lori* knows every student, and he left. This is easily expressed set theoretically. Our domain of entities is $E = \{\text{lori, alex, sara, pim}\}$; let $[[w]]$ indicate the interpretation of w .

This amounts to saying that, for example, the relation *know* is the set of pairs $\langle \alpha, \beta \rangle$ where “ α knows β ”; or that *student* is the set of all those elements which are a student. The case of *every student* is more interesting: it denotes the set of properties that every student has.

$$\begin{aligned}
 (11) \quad [[\text{every student}]] &= \{Y \subseteq E \mid [[\text{student}]] \subseteq Y\} = \{\{\text{alex, sara, lori}\}, \{\text{alex, sara, lori, pim}\}\} \\
 &= \{[[\text{student}]], [[\text{sing}]]\}
 \end{aligned}$$

Alternatively, the words above can be seen as functions from the elements of their set to truth values. For instance, an intransitive verb like *left* can be seen as a function that, given an entity e as argument, will return a truth value t ($\text{left}(\text{lori}) = \text{true}$, $\text{left}(\text{alex}) = \text{false}$). The standard way of representing functions is by means of lambda terms: each word can be assigned a lambda term representing its meaning. For instance, a transitive verb like, *knows* or *trust*, is represented as a function with two arguments (or a function taking the ordered pair of the element in the sets above). To see expressions as functions allows to consider natural language expressions as built by means of function application. Richard Montague introduced the idea that the step that put

together the syntactic structure are associated with the ones that give instructions for assembling meaning representations: the syntactic and semantic rules in his grammar work in parallel. Here, we introduce those concepts of Montague Grammar that are essential for the understanding of our proposal, namely Categorical Grammar and the connection between syntax and semantics.

Syntax-Semantics In Categorical Grammar (CG) syntactic categories can be atomic or complex. The latter are assigned to those expressions that are interpreted as functions and are designed to distinguish whether the argument A must occur to the left (A/B) or the right (B/A) of the function. Hence, function application comes in two versions, a Forward Application (FA) and a Backward Application (BA):

$$(12) \quad (FA) \quad \begin{array}{c} B \\ \swarrow \quad \searrow \\ B/A \quad A \end{array} \quad (BA) \quad \begin{array}{c} B \\ \swarrow \quad \searrow \\ A \quad A \backslash B \end{array}$$

The possibility to express whether the argument must occur on the left or the right of the function is important, for instance, in the case of the two arguments taken by the transitive verb: In English the object occurs on its right and the subject on its left. Having in mind the observation made above, viz. that a category is a name for the set of expressions that belong to it, we can look at function application in more precise terms. Given the categories A , B , C and D such that $C \subseteq A$ and $B \subseteq D$, i.e. every expression in C is also in A and every expression in B is also in D , then if a functor category combines with elements of A as an argument (i.e. it's either $A \backslash \cdot$ or \cdot / A), it also combines with elements of C ; the result of the application is an expression that belongs to the set B and hence to any D of which B is a subset. We will come back to this observation when looking at the dual Lambek Calculus in Section 3.

$$(13) \quad (FA) \quad \begin{array}{c} B \subseteq D \\ \swarrow \quad \searrow \\ B/A \quad C \subseteq A \end{array} \quad (BA) \quad \begin{array}{c} B \subseteq D \\ \swarrow \quad \searrow \\ C \subseteq A \quad A \backslash B \end{array}$$

The whole picture of our little lexicon is summarized in Table 2, where for each word is given the syntactic categorial category, its corresponding semantic type as well as the lambda term of that type that represents the set theoretical meaning of the word given in the last column of the table. We indicate only one representative for each category.^{1,2}


1. Following the same procedure, *every* is assigned the type $(e \rightarrow t) \rightarrow ((e \rightarrow t) \rightarrow t)$ and the term $\lambda Z. \lambda Y. \forall \lambda x. ((Z x) \Rightarrow (Y x))$. An alternative notation is $\lambda Z. \lambda Y. \forall x. ((Z x) \Rightarrow (Y x))$.
2. The mapping from syntactic categories to semantic types, $\text{type} : \text{CAT} \rightarrow \text{TYPE}$, is given below.

$$\begin{array}{ll} \text{type}(np) = e; & \text{type}(A/B) = (\text{type}(B) \rightarrow \text{type}(A)); \\ \text{type}(s) = t; & \text{type}(B \backslash A) = (\text{type}(B) \rightarrow \text{type}(A)); \\ \text{type}(n) = (e \rightarrow t). \end{array}$$

We can now look at an example highlighting the syntax-semantics link captured by CG and its connection with the λ -calculus. The syntactic tree (tree on the left) is built by (FA) and (BA) and each step in it has a one-to-one correspondence with the λ -rules applied to assemble the semantic representation (the tree on the right).

Table 2. Syntactic categories and Semantic types

word	category	type	λ -term	meaning
<i>alex</i>	<i>dp</i>	e	alex	alex
<i>left</i>	$dp \backslash s$	$e \rightarrow t$	$\lambda x_e. (\text{left } x)_t$	$\{\text{lori}\}$
<i>trusts</i>	$(dp \backslash s) / dp$	$e \rightarrow (e \rightarrow t)$	$\lambda x_e. (\lambda y_e. ((\text{trusts } x) y))_t$	$\{\langle \text{alex}, \text{alex} \rangle\}$
<i>student</i>	n	$e \rightarrow t$	$\lambda x_e. (\text{student } x)_t$	$\{\text{alex}, \text{sara}, \text{lori}\}$
<i>every student</i>	$(s / dp) \backslash s$	$(e \rightarrow t) \rightarrow t$	$\lambda Y. \forall \lambda x. ((\text{student } x) \rightarrow (Y x))$	$\{Y \subseteq E \mid [[\text{student}]] \subseteq Y\}$



2.2 The mathematics of sentence structure (Lambek '58 and '61)

Lambek (Lambek 1958) showed that the deduction theorem, that is the backbone of Logic, is also at the heart of linguistic structure composition. With this application in mind, the theorem says that if a linguistic structure Γ composed with another constituent of syntactic category A is of category B then is of category B given A , or alternatively, it is a B missing a A

$$(15) \quad \text{If } A, \Gamma \vdash B \text{ then } \Gamma \vdash A \rightarrow B$$

For instance, if *Lori knows the student* is a sentence then *knows the student* is a sentence missing a determiner phrase –given that *Lori* is of the latter category (*dp*).

$$(16) \quad \text{If } \text{Lori}_{dp}, [\text{knows the student}] \vdash s \text{ then } [\text{knows the student}] \vdash dp \rightarrow s$$

The connection above between the comma and the implication is known in algebra as residuation principle and it's an “iff”. This principle is also behind the more familiar mathematical operators of multiplication (\times) and division ($-$):

$$(17) \quad x \times y \leq z \text{ iff } y \leq \frac{z}{x}$$

The reader can compare the similar behaviour shown by the \times and $-$ with respect to the \leq with the one of the $,$ and \rightarrow with the \vdash . To learn from this comparison, one can observe that the multiplication is characterized by some well known properties like commutativity, associativity and distributivity over the $+$:

$$(18) \quad x \times y = y \times x \quad (x \times y) \times z = x \times (y \times z) \quad (x + y) \times z = (x \times z) + (y \times z)$$

Similarly, we are interested in knowing which are the properties of the “comma” above when the concatenated elements are words. In the approach we follow, the “comma” is taken to be not-commutative and not-associative to avoid over-generation problems. (For other choices, see (Steedman 2000), among others).

As the comma’s behavior varies, so does the conditional’s. In particular, when rejecting commutativity of the comma the implication splits into left implication ($A \backslash B$, “if A on the left then B ”) and right implication (B / A , “ B if A on the right”), since the order of occurrences of the formulas matters. Hence, the residuation principle becomes:

$$(19) \quad (a1) \quad A \otimes C \vdash B \text{ iff } C \vdash A \backslash B \\ (a2) \quad C \otimes A \vdash B \text{ iff } C \vdash B / A$$

The difference with the comma is emphasised by using the \otimes operator in its place.

NL (Non-associative Lambek calculus), introduced in (Lambek 1961), is the pure calculus of residuation since it is characterized by only these two rules (besides transitivity, viz., if $A \vdash B$ and $B \vdash C$, then $A \vdash C$, and reflexivity, viz. $A \vdash A$, of the \vdash).

As we have observed when introducing CG, distinguishing the left and right implication turns out to be relevant when formulas stand for syntactic categories of words. By applying the residuation principle we obtain the category of a verb phrase as well as the one of a transitive verb.

$$(20) \quad (a1) \quad \text{Lori}_{dp} \otimes [\text{knows the student}] \vdash s \text{ iff } [\text{knows the student}] \vdash dp \backslash s \\ (a2) \quad [\text{knows}]_{tv} \otimes [\text{the student}]_{dp} \vdash dp \backslash s \text{ iff } [\text{knows}]_{tv} \vdash (dp \backslash s) / dp$$

Furthermore, the object of our toy example is a dp consisting of a noun (*student*) that is assigned the atomic category n and an article (*the*). Its category is determined in the, by now, usual way:

$$(21) \quad [\text{the}] \otimes [\text{student}]_n \vdash dp \text{ iff } [\text{the}] \vdash dp / n$$

Forward and Backward function applications are theorems of this calculus:

$$(22) \quad B / A \otimes A \vdash B \text{ and } A \otimes A \backslash B \vdash B$$

The Lambek Calculus is the logic of the *merging* relation represented by the \otimes (See (Vermaat 2005) for a detailed comparison of this relation with the Merge relation of the Minimalist Framework):

- (23) If there are two expressions $e_1 \in B/A$, $e_2 \in A$ s.t. $\text{Merge}(e_1, e_2, e_3)$, i.e. e_3 is the result of merging e_1 and e_2 , then by definition of \otimes , $e_3 \in B/A \otimes A$, and, hence $e_3 \in B$, since $B/A \otimes A \vdash B$. An expression e_1 belongs to a functor category B/A if for all expressions e_2, e_3 such that $\text{Merge}(e_1, e_2, e_3)$ and $e_2 \in A$, then $e_3 \in B$.

2.3 “Parsing as deduction”: Sequent Calculus

A logic can be presented in several proof-system styles. Natural Deduction is the most well known one among linguists familiar with Logic. Other systems are Gentzen Sequent Calculus, Tableaux, Proof nets. Here we will use Sequent Calculus since it is the most suitable for introducing the concept of dual proofs and symmetric calculus that are at the heart of our proposal. It is a “decision procedure” and we use it to determine whether a phrase is a well-formed expression of a given category (and to assign it a meaning representation in the process of checking this.)

Gentzen Sequents Sequent Calculus is a proof system introduced by Gentzen to reason on structures. A *sequent* has the form

$$(24) \quad \Sigma \vdash \Gamma$$

where both Σ and Γ are sequences of logical formulas (i.e., both the number and the order of the occurring formula matter). The symbol \vdash is usually referred to as turnstile and is often read, suggestively, as “yields” or “proves”; it is not a symbol in the language, rather it is a symbol in the metalanguage used to discuss proofs. In the sequent above, Σ is called the antecedent and Γ is said to be the succedent of the sequent. A sequent derivation starts from the sequent to be proved and by application of the inference rules reduces it into simpler sequents (roughly, a sequent with a lower number of operators). In other words, inference rules are precise instructions to build the derivation of a given sequent. The choice of which inference rule to apply is dictated by the main operators of the formulas in the sequent. As we will show later, in some cases more choices are possible. A derivation is concluded when axiom links are reached.

The properties of the calculus determines the actual inference rules which are written with a list of sequents above (premises) and below a line (conclusion). Each operator has an inference rule that eliminates it from the left of the turnstile and an inference rule that eliminates it from the right.³ For instance, the logical rule that

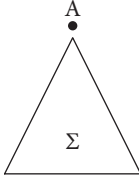
3. In Natural Deduction, the sequent rule introducing an operator on the right corresponds to the Introduction rule, whereas the introduction of the operator on the left of the \vdash corresponds to the Elimination rule.

eliminates \wedge of propositional logic from the left of the turnstile is represented as following

$$(25) \quad \frac{\Gamma[A] \vdash \Delta}{\Gamma[A \wedge B] \vdash \Delta} (\wedge L)$$

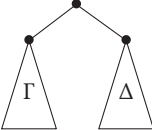
The rule says: the structure containing the formula $A \wedge B$ yields the structure Δ , if Γ containing A can be proved to yield Δ , i.e. the problem of proving the sequent in the conclusion is reduced to a simpler problem by eliminating the \wedge .

Non-associative Lambek Calculus In the case of the Non-associative Lambek calculus (NL) the property encoded into the inference rules is the residuation principle discussed above. Here we do not go into the details of how this principle is captured by the inference rules of the Sequent proof system, and refer the reader interested in the issue to (Areces & Bernardi 2004). Since in NL the sequence can be seen as tree-structure, to help understanding the austere sequent notation for each inference rule of the calculus we give a tree representation as well. In NL the succedent of sequents is not a structure but a formula:

$$(26) \quad \Sigma \vdash A$$


This statement expresses the judgement that the tree structure Σ is of category A . Thinking of categories as labels for sets of expressions, the sequent can be read as a proof of the fact that every expression in Σ is also in A . We will thus take the liberty of saying that the sequent proves that the structure Σ is of category A .

Recall, that trees are bracketed strings of formulas. The smallest tree is just a single formula A . Composite trees are formed by taking two trees Γ and Δ together as immediate daughters of a binary branching structure. In the sequent notation, we write (Γ, Δ) . Graphically,

$$(27)$$


Axiom, Cut: How can we arrive at a judgement $\Gamma \vdash A$? In the case where Γ is just the single node tree consisting of the formula A , we are done. In the language of the sequent calculus, we have the axiom $A \vdash A$ (i.e. it simply expresses the reflexivity of the \vdash). We will graphically represent this case as

$$(28) \quad \begin{array}{c} A \\ \bullet \\ A \end{array}$$

where we write the formula at the left-hand side of \vdash below the node \bullet , and the formula at the right-hand side above it.

The axiom case deals with the simplest possible trees — trees consisting of one single formula. Next we need rules to reason about composite trees. Consider first the situation where we have a tree which has among its leaves a formula B . Suppose we have arrived at the judgment that this tree is of category C . We can graphically depict this situation as

$$(29) \quad \begin{array}{c} C \\ \bullet \\ \Gamma \\ \bullet \\ B \end{array}$$

In the language of sequent calculus we write $\Gamma[B] \vdash C$. The square bracket notation $\Gamma[\Gamma']$ is used for picking out a distinguished sub-tree Γ' in a tree Γ . In this particular case, the sub-tree Γ' is the single-node tree B .

Suppose now that we also have a second tree Δ which we can show to be of category B . In that case, substituting the tree Δ for the node B in Γ will produce a tree that is still of category C . In other words: in a well-formed phrase of category C , we can replace any of its constituents by another constituent, as long as the replacement and the original constituent has the same category.

This grammaticality-preserving substitution yields the inference rule below, which is known as the Cut rule.

$$(30) \quad \begin{array}{ccc} \text{from} & \begin{array}{c} B \\ \bullet \\ \Delta \end{array} & \text{and} & \begin{array}{c} C \\ \bullet \\ \Gamma \\ \bullet \\ B \end{array} & \text{infer} & \begin{array}{c} C \\ \bullet \\ \Gamma \\ \bullet \\ \Delta \end{array} \end{array}$$

In the concise sequent notation, we write⁴

$$(31) \quad \frac{\Delta \vdash B \quad \Gamma[B] \vdash C}{\Gamma[\Delta] \vdash C} \quad (\text{Cut})$$

4. The cut rule involves an unknown formula that disappear from the conclusion (given the problem to solve, $\Gamma[\Delta] \vdash C$, it is not known what is the cut-formula, B .) Hence, for decidable proof search, people check whether in a given system the cut rule is admissible, viz. all the theorems that can be proved with the cut rule, can be proved also without it.

Logical rules: The Axiom and the Cut rule do not mention particular category-forming operations: they hold in full generality for any formula. Let us turn then to the inference rules for the category-forming operations. For each of the connectives of our grammar logic, $/$, \otimes , \backslash , we need a rule that tells us how to deal with it when we find it right of the derivability sign \vdash , or left of it. The first kind of rules are called *rules of proof*: they tell us how to arrive at a judgement $\Gamma \vdash B$ in the case where B is a complex formula: B/A , $A \backslash B$, or $A \otimes B$. The rules for the connectives left of \vdash are called *rules of use*: they tell us how we can put the tree together in accordance with the interpretation of the connectives $/$, \otimes , \backslash .

Recall the substitution rule (Cut) that allows us to replace any leaf formula B by a tree structure that can be shown to be of category B . One possibility of doing that, in a way that introduces a slash formula B/A , would be to replace the formula B by the tree below, putting together B/A with a formula A to its right – the familiar application schema.

$$(32) \quad \begin{array}{c} \bullet \\ \swarrow \quad \searrow \\ \bullet \quad \bullet \\ B/A \quad A \end{array}$$

More generally (and using our substitution reasoning again), we obtain a tree of category B by combining B/A with any sister tree Δ on its right, provided we can show Δ is of category A . Putting these steps together, we obtain the $(/L)$ inference rule, introducing $/$ in the antecedent tree structure.


$$(33) \quad \begin{array}{ccc} \text{from} & \begin{array}{c} A \\ \bullet \\ \triangle \\ \Delta \end{array} & \text{and} \quad \begin{array}{c} C \\ \bullet \\ \triangle \\ \Gamma \\ \bullet \\ B \end{array} & \text{infer} & \begin{array}{c} C \\ \bullet \\ \triangle \\ \Gamma \\ \bullet \\ \swarrow \quad \searrow \\ B/A \quad \begin{array}{c} \bullet \\ \triangle \\ \Delta \end{array} \end{array} \end{array}$$

In sequent notation, we have:

$$(34) \quad \frac{\Delta \vdash A \quad \frac{\frac{A \vdash A \quad B \vdash B}{B/A \otimes A \vdash B} (/L) \quad \Gamma[B] \vdash C}{\Gamma[B/A \otimes A] \vdash C} (\text{Cut})}{\Gamma[B/A \otimes \Delta] \vdash C} (\text{Cut})$$


In other words, the function application rule can be applied within a tree and its compact formulation is as below.

(35)




The elimination of the $/$ from the left of the turnstile, $(/L)$, is the application of a function of category B/A to a structure Δ , i.e. the sub-structure $B/A \otimes \Delta$ is proved to be of category B , if Δ can be proved to be of category A . The application happens within a bigger structure Γ . Therefore, the sequent is reduced into two simpler sequents. In particular, in the premise in the right branch the sub-structure $B/A \otimes \Delta$ is replaced by its category B . In other words, the sequent in the conclusion holds if it can be proved that Δ yields A and the structure Γ containing B yields C . Notice that if Γ is a structure consisting only of B/A and Δ , then the right branch is $B \vdash C$. Similarly for the $A \setminus B$ case.

Consider the sentence *Lori knows Alex*. We have seen that the words in this sentence can be assigned the category dp , $(dp \setminus s)/dp$ (transitive verb) and dp , respectively. We want to show that we obtain a tree of category s if we put together these words with the constituent structure '[Lori [knows Alex]]'. The inference steps that lead to that conclusion are given below. For ease of reference, we number the atomic formulas. The first step is to decompose our original problem into two smaller sub-problems by applying the $(/L)$ rule. We check, in other words, whether the transitive verb $(dp_2 \setminus s_1)/dp_3$ indeed has a tree of category dp as its right sister. The atomic tree dp_4 for the direct object *Alex* qualifies. What remains to be shown is that the subject verb-phrase combination $(dp_1, (dp_2 \setminus s_1))$ is indeed a tree of category s .



This is shown by applying the $(\setminus L)$ rule.



This rather laborious graphical account of the derivation can be written down more compactly in the sequent notation:

$$(38) \quad \frac{dp_4 \vdash dp_3 \quad \frac{dp_1 \vdash dp_2 \quad s_1 \vdash s_2}{(dp_1, (dp_2 \backslash s_1)) \vdash s_2} (\backslash L)}{(\underbrace{dp_1}_{[Lori]}, \underbrace{((dp_2 \backslash s_1)/dp_3)}_{[knows]}, \underbrace{dp_4}_{[Alex]}) \vdash s_2} (/L)$$

The $(/L)$ and $(\backslash L)$ rules provide instructions on how to use formulas B/A , $A \backslash B$ to build tree structures. In sum, by application of inference rules, i.e. by structure simplification, we have reduced the whole structure $dp_1 \otimes ((dp_2 \backslash s_1)/dp_3 \otimes dp_4)$ to s_1 which matches the top formula $s_2 (s_1 \vdash s_2)$. In the tree we represent this match between the s -categories since it will play a role in the description of the QP analysis:

$$(39) \quad \begin{array}{c} s_1 \vdash s_2 \\ \swarrow \quad \downarrow \quad \searrow \\ \underbrace{dp_1}_{Lori} \quad \underbrace{(dp_2 \backslash s_1)/dp_3}_{knows} \quad \underbrace{dp_4}_{Alex} \end{array}$$

Let us turn then to the rules that tell us how we can decide that a tree structure Γ is of category B/A or $A \backslash B$. Given the interpretation of the slashes, this will be the case if Γ together with a sister node A to its right or left, can be shown to be a tree of category B . For formulas B/A , we obtain the $(/R)$ inference rule below.

$$(40) \quad \text{from } \begin{array}{c} B \\ \swarrow \quad \searrow \\ \Gamma \quad A \end{array} \quad \text{infer } \begin{array}{c} B/A \\ \triangle \\ \Gamma \end{array}$$

The sequent notation is given below, together with the $(\backslash R)$ rule.

$$(41) \quad \frac{(\Gamma, A) \vdash B}{\Gamma \vdash B/A} (/R) \quad \frac{(A, \Gamma) \vdash B}{\Gamma \vdash A \backslash B} (\backslash R)$$

An interesting use of the right rules is the theorem $A \vdash C/(A \backslash C)$, an instance of which is the type lifting rule known to linguists, viz. from e to $(e \rightarrow t) \rightarrow t$. Any word of category dp is proved to be of category $s/(dp \backslash s)$. In other words, any expression that is in the category of determiner phrases is also in the category of quantifier phrases; any expression that is interpreted as an entity can be interpreted as the set of its properties (compare Table 2).

$$(42) \quad \begin{array}{c} dp_1 \vdash dp_2 \quad s_2 \vdash s_1 \\ \hline dp_1 \otimes dp_2 \backslash s_2 \vdash s_1 \\ \hline dp_1 \vdash s_1/(dp_2 \backslash s_2) \end{array} \begin{array}{c} (\backslash L) \\ \\ (/R) \end{array} \quad \begin{array}{c} s_1/(dp_2 \backslash s_2) \\ \bullet \\ dp_1 \end{array} \quad \text{If } \begin{array}{c} s_1 \\ \swarrow \quad \searrow \\ dp_1 \quad dp_2 \backslash s_2 \end{array} \quad \text{If } \begin{array}{cc} dp_2 & s_1 \\ \bullet & \text{and } \bullet \\ dp_1 & s_2 \end{array}$$

Since in the sequent in the conclusion there is only one main operator, $/$, and it occurs on the right of the turnstile, the only rule that can be applied is $(/R)$, hence the sequent is reduced to a simpler sequent (a sequent with less implication operators), viz. $dp_1 \otimes dp_2 \backslash s_2 \vdash s_1$. This sequent holds if $dp_1 \vdash dp_2$ and $s_2 \vdash s_1$, since both are axioms we are done. Notice that $s/(dp \backslash s) \not\vdash dp$, as one would expected by thinking of the domains of interpretation of the corresponding functions, viz. it is not true that whatever can be interpreted as a set of properties can be interpreted as an entity.

This theorem is an essential part of the proofs of “Argument Lowering”, $B/(C/(A \backslash C)) \vdash B/A$, and “Value Raising”, $B/A \vdash (C/(B \backslash C))/A$, used by Hendriks (Hendriks 1993) to account for scope ambiguities. We give the sequent derivation, the reader is invited to translate them into a tree notation.

$$(43) \quad \frac{\frac{\frac{\vdots}{A \vdash C/(A \backslash C)} \quad B \vdash B}{B/(C/(A \backslash C)) \otimes A \vdash B} (/L) \quad \frac{\frac{\vdots}{A \vdash A} \quad B \vdash C/(B \backslash C)}{B/A \otimes A \vdash (C/(B \backslash C))} (/L)}{\frac{B/(C/(A \backslash C)) \vdash B/A}{B/A \vdash (C/(B \backslash C))/A} (/R)} (/R)$$

Hendriks uses also a third rule called “Argument Raising”: this rule is not a theorem of the Logic, since as we have mentioned above $C/(A \backslash C)$ does not derive A and hence the proof fails to reach all axioms.⁵ (See (Bastenhof 2007) for a detailed discussion on Argument Raising.)

$$(44) \quad \text{FAILED} \quad \frac{\frac{\frac{\vdots}{C/(A \backslash C) \vdash A} \quad A \vdash A}{B/A \otimes C/(A \backslash C) \vdash B} (/L) \quad \frac{\vdots}{B/A \vdash B/(C/(A \backslash C))} (/R)}{B/A \vdash B/(C/(A \backslash C))} (/R)$$

Bottom-up vs. Top-down parsing strategies Let us now consider the sentence *The student who knows Sara left*. The relative pronoun *who* works as a bridge between a sentence missing a dp in subject position, *knows Sara*, and the noun *student* that the relative clause modifies, viz. $who \in (n \backslash n)/(dp \backslash s)$ as we can compute by means of the residuation principle:


$$(45) \quad [who]_{rel} \otimes [knows Sara]_{dp \backslash s} \vdash n \backslash n \text{ iff } [who]_{rel} \vdash (n \backslash n)/(dp \backslash s)$$

This example will help us bringing in another important observation regarding parsing strategies. In the examples seen so far, at each step there was only one possible

5. This failure is due to the monotonicity property of the implication, that is downward monotonic in the argument position and upward monotonic in the value position, viz. $\neg \backslash +$ and $+/-$.

inference rule that could have been applied. However, it can happen that a sequent could be proved in more than one way. For instance, it contains more than one complex formula either in the antecedent or the succedent position. One can choose to follow either a (i) bottom-up or a (ii) top-down strategy as we will illustrate with the example below.

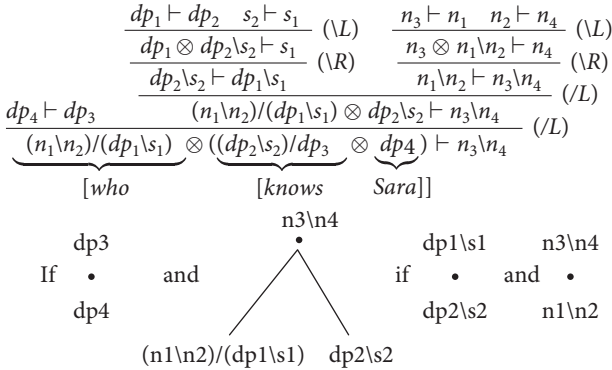
The problem to solve is to check whether the tree below is of category $(n \setminus n)$.



At each inference step, we follow a bottom-up strategies: We start from the bottom of the tree. In the leaves, there are two complex formulas that could be activated, hence there are two possible inference rules to apply as a first step: we could simplify the sequent either (i) by applying the function $(dp_2 \setminus s_2) / dp_3$ to dp_4 or (ii) by applying the function $(n \setminus n) / (dp \setminus s)$ to the sub-structure $((dp_2 \setminus s) / dp_3 \otimes dp_4)$. In a bottom-up approach the choice is guided by the atomic categories in the leaves, hence, in this case by dp_4 . The function to activate first is, therefore, the one of the transitive verb whose argument matches the atomic leave. The other inference steps are applied using these same criteria. In the tree notation, we leave the last two trees unfolded. The reader can check by herself the last steps that lead to the axiom links indicated in the sequent.

(47)


Bottom-up strategy:



The top-down strategy is illustrated below both in the sequent and in the tree notation. The starting point is the top-formula $n_3 \setminus n_4$, that can be simplified by $(\setminus R)$. The next step is the inference rule whose result matches the new top formula, i.e. n_4 . Hence, the function activated first is the category of *who*. Again we leave to the reader the unfolding of the last trees into the axiom links.

(48)

$$\frac{\frac{\frac{dp_1 \vdash dp_2 \quad s_2 \vdash s_1}{dp_1 \otimes dp_2 \vdash s_2 \vdash s_1} (\wedge L) \quad \frac{dp_4 \vdash dp_3 \quad dp_2 \vdash s_2 \vdash dp_1 \vdash s_1}{(dp_2 \vdash s_2)/dp_3 \otimes dp_4 \vdash dp_1 \vdash s_1} (\wedge R)}{\frac{dp_4 \vdash dp_3}{(dp_2 \vdash s_2)/dp_3 \otimes dp_4 \vdash dp_1 \vdash s_1} (\wedge L) \quad \frac{n_3 \vdash n_1 \quad n_2 \vdash n_4}{n_3 \otimes n_1 \vdash n_2 \vdash n_4} (\wedge R)}{\frac{(dp_2 \vdash s_2)/dp_3 \otimes dp_4 \vdash dp_1 \vdash s_1}{n_3 \otimes ((n_1 \vdash n_2)/(dp_1 \vdash s_1)) \otimes ((dp_2 \vdash s_2)/dp_3 \otimes dp_4) \vdash n_4} (\wedge L)}{\frac{(n_1 \vdash n_2)/(dp_1 \vdash s_1) \otimes ((dp_2 \vdash s_2)/dp_3 \otimes dp_4) \vdash n_3 \vdash n_4}{\underbrace{(n_1 \vdash n_2)/(dp_1 \vdash s_1)}_{[who]} \otimes \underbrace{((dp_2 \vdash s_2)/dp_3 \otimes dp_4)}_{[knows \quad Sara]} \vdash n_3 \vdash n_4} (\wedge R)}$$




The first toy example shows that NL captures the composition principle involved in kernel sentences – sentences consisting of local dependencies only (dependencies among juxtaposed constituents). It also illustrates how the concepts of dependency and constituent, so important when looking at natural language, are encoded in the derivation. The dependency between the verb and its arguments are captured by the axiom links involving the determiner phrases, on the other hand the brackets put around the structure at each step in the derivation build the constituents out of the single words. The last example shows that NL also properly handles extraction when it happens from a peripheral (accessible) position.

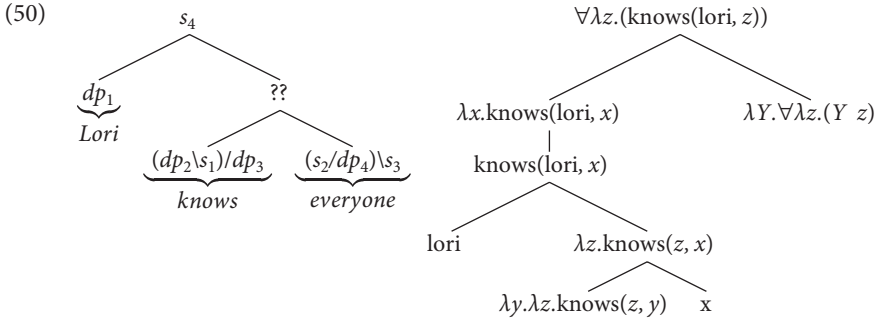
2.4 Lambek calculus limitations

In NL function application can happen only on juxtaposed structures, abstraction only from peripheral position, and there is no way of re-bracketing the tree once built. As a consequence, the system fails to recognize long distance dependencies and cross-dependencies, and does not capture the behavior of in-situ binders, like natural language quantifiers. Let's look at their case now.

An important concept to introduce regards the notion of “scope” of quantifiers. Like in logic, the scope of a quantifier is the part of the expression on which it performs its action; the quantifier binds the variable x of category dp in its scope, e.g. the action of \forall below is to check all the possible assignments of values to the variable z it binds.




The category that we have shown to be derivable from dp and that is mapped into the semantic type $(e \rightarrow t) \rightarrow t$ works properly in NL when in subject position. However, this category does not capture the behaviour of QP in general, e.g. it already causes a failure of the proof when assigned to a QP in object position, as marked by the question marks in the syntactic tree.



By comparing the relation between the syntax and the semantic trees of the first example with the one of this second example, one can see that quantifier may cause a mismatch between the two trees (in the syntactic tree the QP is in an embedded position, whereas in the semantic tree it is in a peripheral position.)

A solution would be to enrich CG with a way to reproduce in the syntactic tree the abstraction used in the semantic one. This solution, however, would cause the generation of a syntactic tree that does not reflect the syntactic constituents, i.e. we would obtain $[[\text{Lori knows}] [\text{every student}]]$. (See (Steedman 2000) for interesting discussion and alternative view on the role played by syntactic constituents). This ability of natural language quantifier to take scope over the structure in which they occur, even if in an embedded position, challenges any formal account of the syntax-semantic interface of natural languages. NL is not expressive enough to account for this behaviour of QP. Our goal is to understand which are the formal properties that characterize this behaviour.

If one wants to mimic the semantic tree at the syntactic level, she could assign to QPs a second category, $(s/dp)\backslash s$ that is also mapped into $(e \rightarrow t) \rightarrow t$ and would need to be able to obtain the following tree on the left, which by function application is reduced to the two simpler trees on the right:



This can be proved by abstraction

$$(52) \quad \begin{array}{c} s_2/dp_4 \\ \swarrow \quad \searrow \\ \underbrace{dp_1}_{Lori} \quad \underbrace{(dp_2 \backslash s_1)/dp_3}_{knows} \end{array} \quad \text{If} \quad \begin{array}{c} s_2 \\ \swarrow \quad \searrow \\ \underbrace{dp_1}_{Lori} \quad \underbrace{(dp_2 \backslash s_1)/dp_3}_{knows} \quad dp_4 \end{array}$$

We have arrived to a tree that cannot be simplified further by means of the Lambek calculus' inference rules. One would first need to re-bracket it,

$$(53) \quad \begin{array}{c} s_2 \\ \swarrow \quad \searrow \\ \underbrace{dp_1}_{Lori} \quad \underbrace{(dp_2 \backslash s_1)/dp_3}_{knows} \quad dp_4 \end{array}$$

The tree obtained can then be simplified by function applications, arriving at the following matches of atomic formulas:

$$(54) \quad \begin{array}{ccc} dp_2 & dp_3 & s_2 \\ \bullet & \bullet & \bullet \\ dp_1 & dp_4 & s_1 \end{array}$$

The re-bracketing corresponds to applying the associativity rule as shown in the sequent below where the corresponding inference step is marked by (Ass); the reader can check the correspondence of each inference rule with the above trees.


$$(55) \quad \begin{array}{c} dp_1 \vdash dp_2 \quad dp_4 \vdash dp_3 \quad s_1 \vdash s_2 \\ \vdots \\ D \\ \vdots \\ \frac{dp_1 \otimes ((dp_2 \backslash s_1)/dp_3 \otimes dp_4) \vdash s_2}{(dp_1 \otimes (dp_2 \backslash s_1)/dp_3) \otimes dp_4 \vdash s_2} \text{ (Ass)} \\ \hline \frac{(dp_1 \otimes (dp_2 \backslash s_1)/dp_3) \otimes dp_4 \vdash s_2}{dp_1 \otimes (dp_2 \backslash s_1)/dp_2 \vdash s_2/dp_4} \text{ (/R)} \\ \hline \frac{dp_1 \otimes (dp_2 \backslash s_1)/dp_2 \vdash s_2/dp_4 \quad s_3 \vdash s_4}{\underbrace{(dp_1 \otimes ((dp_2 \backslash s_1)/dp_2 \otimes (s_2/dp_4) \backslash s_3))}_{[Lori] \quad [knows] \quad [everyone]]} \vdash s_4 \text{ (\backslash L)} \end{array}$$

The derivation above proves that the structure [Lori [knows everyone]] belongs to s_4 . However, adding associativity (Ass) to NL, i.e. going back to the Lambek's '58 paper, and using two different categories for quantifier in subject and object position is still not a general solution: the QP won't be able to jump out from any position, but only from those that can become peripheral by associativity, e.g. [Lori [thinks [Alice read a book] while she was ill]] would receive only one reading.

Still it is interesting to reflect on the two examples with QP in subject and object position to grasp proof-theoretical insights on the syntax-semantic interface of natural language structures involving scope operators.

Syntax-Semantics seen from the axiom links The above example suggests that we can read the scope of the functions in the parsed structure out of the axiom links of the s -atomic formulas: the fact that *Everyone* has scope over the whole-structure (has wide scope) is captured by the axiom link $s_3 \vdash s_4$; the fact that the transitive verb is in the immediate scope of the quantifier is captured by the the axiom link $s_1 \vdash s_2$. Furthermore, the fact that the quantifier binds the variable taken as object by the transitive verb is captured by the axiom link $dp_4 \vdash dp_3$.

We can elaborate further what we have just observed by saying that a quantifier could intuitively be considered as consisting of two components tied to each other: a dp and a function that takes the sentential domain to produce a sentence. The syntactic component (dp) should stay in the embedded position where it works as argument of a verb, and the scope component should be free to travel through the tree so to reach its top level, without modifying the constituent structure. Below we focus on the scope component and represent it schematically as s_{val}/s_{arg} (compare $(s_{arg}/dp)\backslash s_{val}$ and $s_{val}/(dp\backslash s_{arg})$).⁶ Generalizing, we consider a structure with two scope-operators (OP_1 , OP_2) and give the two possible scoping order possibilities.



The trees are consistent with the standard idea that the later a quantifier is introduced in the semantic structure the wider its scope. In Section 4, we discuss the formal properties that characterize the function s_{val}/s_{arg} and its connection with the dp component as well as the composition relation holding between it and the structure over which it takes scope.

In short, the QP category that can be assigned using residuated operators gives rise to the following problems: (i) it requires multiple category assignments for the QP in subject and object position; it fails to produce both (ii) local and (iii) non-local

6. In (Bernardi & Szabolcsi 2008) this observation is exploited to account for the different scope behaviours of quantifiers. The argument- s determines what the operator can immediately scope over. The value- s determines what the operator can be in the immediate scope of.

scope ambiguity. We show that using dual residuated operators, QP can be assigned one category that overcomes these three problems. Reasoning with natural language structure requires some categories to *jump out* of the structure where they sit leaving, however, the structure of the other components unchanged. We show that this ability is captured by the calculus consisting of residuation, dual residuation and the distributivity principles.

3. “The mathematics of sentence structure” revised

In this section we introduce the concept of dual rules and dual proofs by looking at the dual of the residuation principle. We show how the jump required by the quantifiers is obtained by means of the distributivity principles.

The reader may recall the De Morgan’s laws, i.e. rules relating pairs of dual logical operators in a systematic manner expressed in terms of negation. For instance,

$$(57) \quad \neg(A \wedge B) = \neg A \vee \neg B$$

which in words says: Since it is false that two things together are true, at least one of them must be false. Saying that this equality holds for propositional logic means that $\neg(A \wedge B) \vdash \neg A \vee \neg B$ and $\neg A \vee \neg B \vdash \neg(A \wedge B)$ are theorems of the Logic.

In (Schröder 1980) it is introduced the concept of dual rules, for instance, the two rules below are duals. Their duality is expressed in terms of the sequent-turnstile, the two rules are perfectly symmetric with respect to the turnstile:

$$(58) \quad \frac{\Gamma[A] \vdash \Delta}{\Gamma[A \wedge B] \vdash \Delta} (\wedge L) \quad \text{dually} \quad \frac{\Gamma \vdash \Delta[A]}{\Gamma \vdash \Delta[B \vee A]} (\vee R).$$

Similarly, one can think of dual theorems. For instance, one could express either the distributivity of the \wedge over the \vee or dually the one of the \vee over the \wedge .

$$(59) \quad \begin{array}{l} A \wedge (B \vee C) \vdash (A \wedge B) \vee (A \wedge C) \quad \text{dually} \quad (C \vee A) \wedge (B \vee A) \vdash (C \wedge B) \vee A; \\ (A \wedge B) \vee (A \wedge C) \vdash A \wedge (B \vee C) \quad \text{dually} \quad (C \wedge B) \vee A \vdash (C \vee A) \wedge (B \vee A). \end{array}$$

\neg is one of the logical operators of propositional logic, whereas it is not part of the logical language of the Lambek calculus. Hence, we cannot define duality in De Morgan’s style at the object-language level. However, we can still define it at the meta-level in terms of the behaviour of the operators w.r.t. the turnstile.

Residuation and Dual Residuation The logical operators of the Lambek Calculus are $(\backslash, \otimes, /)$; we introduce new operators $(\oslash, \oplus, \oslash)$ and define them as duals of the former in terms of their behaviour with respect to the \vdash . The dual residuation principle was studied by V.N. Grishin (Grishin 1983) and further investigated by Lambek in (Lambek 1993) and Rajeev Goré in (Goré 1997).

We remind the residuation principle by copying it in (a1) and (a2). Their duals are in (b1) and (b2) respectively.

$$(60) \quad \begin{array}{ll} (a1) & A \otimes C \vdash B \text{ iff } C \vdash A \backslash B \quad \text{and} \quad (a2) \quad C \otimes A \vdash B \text{ iff } C \vdash B/A \\ (b1) & B \vdash C \oplus A \text{ iff } B \oslash A \vdash C \quad \text{and} \quad (b2) \quad B \vdash A \oplus C \text{ iff } A \oslash B \vdash C \end{array}$$

In Section 2.2, while introducing the merge relation, \otimes , and the directional implications, \backslash , $/$, we saw that multiplication and fraction obey the residuation principle too. Similarly, there are familiar mathematical operations that obey the dual residuation principle, namely addition, $+$, and difference $-$.

$$(61) \quad y \leq z + x \text{ iff } y - x \leq z$$

We also discussed that while \times enjoys associativity and commutativity, \otimes does not and the lack of the latter causes having two functional implications \backslash and $/$. Similarly, while $+$ has these two properties \oplus does not and hence two “difference” operators exist \oslash and \ominus .

Let us pause on the new type of functions we have introduced with the Dual calculus and the new relation among expressions. We have the dual function (or co-function) application:

$$(62) \quad \begin{array}{lll} B/A \otimes C \vdash D & \text{if} & C \vdash A \quad \text{and} \quad B \vdash D \\ \text{dually} & & \text{dually} \quad \text{dually} \\ D \vdash C \oplus (A \oslash B) & \text{if} & A \vdash C \quad \text{and} \quad D \vdash B \end{array}$$

The backward application and the dual backward application are in the same duality relation.

As we said, the Lambek calculus is the logic for *merging* two expressions into a new one; dually the Dual Lambek calculus is the logic for merging contexts into a new one. Hence, we will call the new relation Merge^c .

$$(63) \quad \text{If } c_3 \in B, \text{ then } c_3 \in A \oplus (A \oslash B), \text{ since } B \vdash A \oplus (A \oslash B); \text{ then, by definition of } \oplus, \\ \text{for all contexts } c_1, c_2 \text{ if } \text{Merge}^c(c_1, c_2, c_3), \text{ viz. } c_3 \text{ is the context resulting by the merge} \\ \text{of the contexts } c_1 \text{ and } c_2, \text{ then } c_1 \in A \text{ or } c_2 \in A \oslash B. \text{ A context } c_2 \text{ belongs} \\ \text{to a co-function category } A \oslash B \text{ if there are two contexts } c_1 \text{ and } c_3 \text{ such that} \\ \text{Merge}^c(c_1, c_2, c_3), c_1 \notin A \text{ and } c_3 \in B.$$

Taking the dual perspective is as if we look at structures in a mirror and hence left and right are swapped around. For instance, the example used to explain residuation in Section 2.2, repeated in Figure 2 as (a), would be seen as (b).


$$(a) \quad \underbrace{dp}_{[\text{Lori}]} \otimes \underbrace{((dp \backslash s_1)/dp \otimes (dp/n \otimes n))}_{[\text{knows } [\text{the student}]]} \vdash s_2 \quad \text{dually} \quad (b) \quad s_2 \vdash \underbrace{((n \oplus (n \oslash dp)) \oplus ((dp \oslash s_1) \oslash dp))}_{[[\text{student the}]\text{ knows}]} \oplus \underbrace{dp}_{[\text{Lori}]}$$

Figure 2. Residuation in a dual perspective


(a) by means of function applications the whole structure in the antecedent is reduced to s_1 which matches the top formula s_2 . Dually (b) by means of dual function

applications the whole structure on the succedent is reduced to s_1 that is matched by s_2 . Notice that in this view, we are reading the entailment in (b) right-to-left, from the succedent to the antecedent. Another way to say this is that both in (a) and in (b) the “focus” is on s_2 that is built out of the other formulas. When necessary for the understanding of the explanation we will mark the left-to-right and right-to-left reading of the entailment as \vdash_{\triangleright} and \vdash_{\triangleleft} , respectively, where the triangle points towards the focused formula.


To emphasize the difference between the structure connected by \otimes and the one connected by \oplus , in the tree notation we will represent the latter as upside down trees. The two trees could also be thought as the result of merging structures into a new one and the expansion of a structures into parts, respectively.



As shown by this example, the sequents of the Lambek calculus have a structure on the succedent, and a formula on the antecedent: $\Sigma \vdash A$; dually, the sequents of the Dual Lambek Calculus have the structure on the antecedent and a logical formula on the succedent of the turnstile.



Lambek Grishin By taking together the Lambek Calculus and its Dual, we obtain sequents of the following shape:



Since the structure in the antecedent, Σ , consists of formulas connected by \otimes we call them \otimes -structure. Similarly, we call the structure in the succedent, Δ , \oplus -structure.

The two logics communicate via the Cut rule that creates an unfocused entailment as illustrated by the example below. Observe that,

$$(67) \quad (B/A) \otimes A \vdash C \oplus (C \otimes B)$$

for every structure if it belongs to $(B/A) \otimes A$, then it belongs to $C \oplus (C \otimes B)$ too. Step-wise can be explained via the communication established between the two merging relations by the Cut rule:

$$(68) \quad \frac{(B/A) \otimes A \vdash_{\triangleright} B \quad B \vdash_{\triangleleft} C \oplus (C \otimes B)}{(B/A) \otimes A \vdash C \oplus (C \otimes B)} \text{ (Cut)}$$

The left premise is read left-to-right (\triangleright) it's governed by the relation merging two expressions into an expressions of category B whereas the right premise is read right-to-left (\triangleleft) it's governed by the relation that merges contexts into a new one of category B^c , hence we reach a tree of category B and a tree with an hole of category B , the former can hence be plugged into the latter reaching and unfocused entailment (marked by the absence of the triangle on the turnstile). In the sequel, we will omit the focus label since we won't go into any proof theoretical details.

The Sequent style function application and co-function application of LG are below. As the reader can see the only difference with respect to the one discussed in Section 2.2 is the presence of a structure in the succedent position, as we have just explained. The $(\otimes R)$ rule is just the dual of $(/L)$.

$$(69) \quad \frac{\Delta \vdash A \quad \Gamma[B] \vdash \Sigma}{\Gamma[B/A \otimes \Delta] \vdash \Sigma} (/L) \quad \text{dually} \quad \frac{A \vdash \Delta \quad \Sigma \vdash \Gamma[B]}{\Sigma \vdash \Gamma[\Delta \oplus (A \otimes B)]} (\otimes R)$$

Still, neither the Lambek calculus nor the dual Lambek calculus allows a category to jump out of an embedded position. What is still missing are the distributivity properties of the \otimes and \oplus over the co-implications (\oslash , \otimes) and the implications (\backslash , $/$), respectively. These properties were first studied by Grishin (Grishin 1983). We call the system consisting of the residuation and dual residuation principles and the distributivity principles below, LG, Lambek Grishin. The distributivity rules are a form of mixed-associativity (MA) and mixed-commutativity (MC) since the former leaves the formulas in their position while changing the brackets, and the latter commutes the formulas while leaving the brackets unchanged.⁷ The term Mix stands for the fact that each principle involves operators from the two families: the residuated triple (\backslash , \otimes , $/$) and its dual (\oslash , \oplus , \otimes). Below we give the principles for \otimes and $/$, similar principles, modulo directionality, hold for their symmetric operators \oslash and \backslash .

7. Note that the distributivity we are used to in mathematics, is the one of \times over the $+$, $(A + B) \times C = (A \times C) + (B \times C)$. It involves resource duplication (see the number of occurrence of C differs in the two sides of the equation), whereas we are working in a resource sensitive system hence no formula is duplicated.

$$(70) \quad (\otimes MA) (B \otimes C) \otimes A \vdash B \otimes (C \otimes A) \text{ dually } (\oplus MA) (A \oplus C)/B \vdash A \oplus (C/B) \\ (\otimes MC) A \otimes (B \otimes C) \vdash B \otimes (A \otimes C) \text{ dually } (\oplus MC) (C \oplus A)/B \vdash (C/B) \oplus A$$

These properties establish how the implication \backslash and $/$ behave when occurring within a \otimes -structure, and similarly how the dual implications \otimes and \oplus behave when occurring within a \otimes -structure. They are structure preserving, in the sense that they respect the non-commutativity and non-associativity of the operations they combine.

In the Sequent system, we compile the distributivity principles into the abstraction and co-abstraction inference rules.⁸

$$(71) \quad \frac{\Gamma \otimes A \vdash \Delta[B]}{\Gamma \vdash \Delta[B/A]} (/R) \quad \text{dually} \quad \frac{\Delta[B] \vdash A \oplus \Gamma}{\Delta[A \otimes B] \vdash \Gamma} (\otimes L)$$

Let's compare the new abstraction rules with the one given in Section 2.2: thanks to the compiled distributivity principles, abstraction can happen within a structure Δ , which is a \oplus -structure. Similarly, co-abstraction happens within a \otimes -structure, ie. the one where a quantifier is sitting in and needs to jump out from. It becomes clear already from the sequent rules how a category can now jump out from an embedded position, but we are going to look them at work in the next section.

The tree representations given in Section 2 need to be modified so to take the structure in the succedent position into account. We give the one of the abstraction rule $(/R)$ that is the most interesting to be considered:

$$(72) \quad \text{from} \quad \begin{array}{c} B \\ \bullet \\ \Delta \\ \bullet \\ \Gamma \quad A \end{array} \quad \text{infer} \quad \begin{array}{c} B/A \\ \bullet \\ \Delta \\ \bullet \\ \Gamma \end{array}$$

8. Below we unfold the steps behind the compiled logical rules. To help understanding the steps, we take Δ to consist of only two \oplus in the $(/R)$ case, and Γ of only two \otimes in the $(\otimes L)$ case. But the two rules apply for arbitrarily big structures.

$$\begin{array}{l} \frac{\Gamma \otimes B \vdash (D \oplus C) \oplus A}{\Gamma \vdash (D \oplus C) \oplus A/B} (RES) \\ \frac{\Gamma \vdash (D \oplus C) \oplus A/B}{\Gamma \vdash ((D \oplus C)/B) \oplus A} (\oplus MC) \\ \frac{\Gamma \vdash ((D \oplus C)/B) \oplus A}{\Gamma \vdash (D \oplus (C/B)) \oplus A} (\oplus MA) \end{array} \quad \frac{\Gamma \otimes B \vdash \Delta[C]}{\Gamma \vdash \Delta[C/B]} (/R)$$

$$\begin{array}{l} \frac{A \otimes (C \otimes D) \vdash B \oplus \Delta}{B \otimes (A \otimes (C \otimes D)) \vdash \Delta} (DRES) \\ \frac{B \otimes (A \otimes (C \otimes D)) \vdash \Delta}{A \otimes (B \otimes (C \otimes D)) \vdash \Delta} (\otimes MC) \\ \frac{A \otimes (B \otimes (C \otimes D)) \vdash \Delta}{A \otimes ((B \otimes C) \otimes D) \vdash \Delta} (\otimes MA) \end{array} \quad \frac{\Gamma[C] \vdash B \oplus \Delta}{\Gamma[B \otimes C] \vdash \Delta} (\otimes L)$$

Figure 3 summarizes the Lambek Grishin calculus in the sequent notation which encodes the mathematical principles represented in Figure 4 that we claim characterize Natural Language Structure. In Figure 4, the symbol \sim marks that the principles of mix-associativity and mix-commutativity are the symmetric version of those without it, i.e. they are about \backslash and \oslash instead of $/$ and \odot respectively. As mentioned above, the axiom stands for the reflexivity (REF) of the \vdash . The cut rule encodes the transitivity (TRAN) of the \vdash in the sequent notation by taking into account the fact that the formulas B can occur within structures (Γ' and Δ') and A and C can be structures, i.e. Δ and Γ , respectively. (See (Bernardi and Moortgat to appear) for the precise sequent derivation, here we overlook some proof theoretical issues.) The abstractions ($/R$, $\backslash R$) and co-abstractions ($\oslash L$, $\odot L$) encode one side of the residuation and dual residuation rule; the other side of residuation is compiled into the function and co-function applications, ($/L$, $\backslash L$) and ($\oslash R$, $\odot R$). The reader is referred to (Areces & Bernardi 2004) for a detailed explanation of the relation between the encoding of the mathematical principles into the sequent system.

$$\begin{array}{c}
A \vdash A \\
\\
\frac{\Delta \vdash \Gamma'[B] \quad B \vdash \Gamma}{\Delta \vdash \Gamma'[\Gamma]} (Cut1) \quad \frac{\Delta \vdash B \quad \Delta'[B] \vdash \Gamma}{\Delta'[\Delta] \vdash \Gamma} (Cut2) \\
\\
\frac{\Delta \vdash A \quad \Gamma[B] \vdash \Sigma}{\Gamma[B/A \otimes \Delta] \vdash \Sigma} (/L) \quad \frac{A \vdash \Delta \quad \Sigma \vdash \Gamma[B]}{\Sigma \vdash \Gamma[\Delta \oplus A \odot B]} (\odot R) \\
\\
\frac{\Gamma \otimes A \vdash \Delta[B]}{\Gamma \vdash \Delta[B/A]} (/R) \quad \frac{\Delta[B] \vdash A \oplus \Gamma}{\Delta[A \odot B] \vdash \Gamma} (\odot L) \\
\\
\frac{\Delta \vdash A \quad \Gamma[B] \vdash \Sigma}{\Gamma[\Delta \otimes A \backslash B] \vdash \Sigma} (\backslash L) \quad \frac{A \vdash \Delta \quad \Sigma \vdash \Gamma[B]}{\Sigma \vdash \Gamma[(B \oslash A) \oplus \Delta]} (\oslash R) \\
\\
\frac{A \otimes \Gamma \vdash \Delta[B]}{\Gamma \vdash \Delta[A \backslash B]} (\backslash R) \quad \frac{\Delta[B] \vdash \Gamma \oplus A}{\Delta[(B \oslash A)] \vdash \Gamma} (\oslash L)
\end{array}$$

Figure 3. LG: Sequent Calculus

$$\begin{array}{c}
(REF) A \vdash A \\
\\
(TRAN) \text{ if } A \vdash B \text{ and } B \vdash C, \text{ then } A \vdash C \\
\\
(RES) C \vdash A \backslash B \text{ iff } A \otimes C \vdash B \text{ iff } A \vdash B/C \\
\\
(DRES) (B \oslash A) \vdash C \text{ iff } B \vdash C \oplus A \text{ iff } C \oslash B \vdash A \\
\\
(\otimes MA) (B \oslash C) \otimes A \vdash B \oslash (C \otimes A) \quad (\oplus MA) (A \oplus C)/B \vdash A \oplus (C/B) \\
(\otimes MC) A \otimes (B \oslash C) \vdash B \oslash (A \otimes C) \quad (\oplus MC) (C \oplus A)/B \vdash (C/B) \oplus A \\
(\otimes MA^\sim) A \otimes (C \oslash B) \vdash (A \oslash C) \oslash B \quad (\oplus MA^\sim) B \backslash (C \oplus A) \vdash (B \backslash C) \oplus A \\
(\otimes MC^\sim) (C \oslash B) \otimes A \vdash (C \otimes A) \oslash B \quad (\oplus MC^\sim) B \backslash (A \oplus C) \vdash A \oplus (B \backslash C)
\end{array}$$

Figure 4. LG: properties

4. Case study: Quantifier phrases

As we have seen in (50) of Section 2, a quantificational expression semantically behaves as if it appeared in a different position than its actual position in the sentence: it exhibits “inverse scope effects”. Because of this ability of jumping out from an embedded position, QPs cause scope ambiguity: the same syntactic structure can receive more than one interpretation. A classical example showing this effect is a sentence with quantifiers as subject and object; it may have two readings if either the subject has scope over the object (1a) or the other way around (1b), though the syntactic structure is shared (in both cases, *someone* is the subject and *everyone* is the object that together with the transitive verb forms the verb phrase). Similarly, scope ambiguity arises when a QP interacts with other scope operators like negation, intentional verbs, wh-phrases in questions, adverbs and coordination. This ability of the quantifiers to scope over the whole structure in which they are embedded can be unbounded, as illustrated by the case of the quantifiers occurring in the complement sentence in (2) of (73). Again, the two interpretations differ in the scope relation (>) between *think* and *every man*, while the syntactic structure is shared.

Natural language offers also cases of bounded scope: structures that delimit the scope of the operators that cannot express their action outside those boundaries, as illustrated by the scope possibility in (3a) that is not suitable for (3) Finally, it has been shown (Szabolcsi 1997) that quantifiers differ in their scope behaviour as exemplified in (4).

- (73) 1. [Someone [knows everyone]_{vp}]_s
 a. There exists someone that everyone knows [someone > everyone]
 b. For everyone there is someone that he knows [everyone > someone]
2. [Lori_{dp} [thinks [every man is immortal]_s]_{vp}]_s
 a. Lori thinks that being immortal is characteristic of men. [thinks > every man]
 b. Lori thinks of every actual man that he is immortal. [every man > think]
3. Two politicians spy on [someone from every city]
 a. *every city > two politicians > someone
4. Lori [didn't [read QP]]
 a. Lori didn't read a book [Not > A], [A > Not]
 b. Lori didn't read every book [Not > Every], [*Every > Not]


As clearly explained in (Szabolcsi 2008), this mismatch between syntax and semantics exhibited by quantifiers might call for a distinction between syntactic domain and semantic scope, where the former has been defined in terms of c-command, maximal projection, feature inheritance, or other similar notions, while the latter is the part of the structure on which the operator performs its action, its scope. The hypothesis

below, originally proposed in (Reinhart 1979), has found many supporters among linguists. Basically, the claim is that syntactic structure “determines” semantic scope.

- (74) “The scope of a linguistic operator coincides with its domain in some syntactic representation that the operator is part of.”

The reader interested in an up-to-date overview of the QP problem and to the several solutions proposed in the literature is referred to (Szabolcsi 2008) and (Ruys and Winter to appear). Here we mention the two well known solutions of the quantifiers puzzle proposed in (May 1977) and in (Hendriks 1993), since they will help highlighting some important aspects of our proposal.

May’s approach produces semantic constituent structures in abstract level. May proposes that syntax does not end with producing the surface string. Instead, a movement rule, called “Quantifier raising” (QR) continues to operate at an abstract level, called the “Logical Form” (LF), and attach each phrase containing a quantifier to its domain by splitting the node S of the surface structure into two nodes. The rule leaves a virtual trace t coindexed with the moved QP. Its earliest formulation operates as shown below.



QR may derive from one given surface structure several different LFs with different scope relations. The structures at the LF level are then interpreted obtaining the meaning representations.

In the solution we propose the syntactic tree is not re-written into a semantic tree, rather the QP category splits into two parts: the syntactic-component, dp , that actually stays in the syntactic structure – where in May’s LF there is a virtual trace t – and it’s only the scope-component that moves out and reaches the top S-node of the sentence where the QP performs its action.

Hendriks dissociates scope from pure syntax in that it allows one to maintain whatever constituent structure seems independently motivated and still delivers all imaginable scope relations: all the possible scope relations can be obtained by means of three type-change rules: “Argument Raising”, “Value Raising”, and “Argument Lowering”. As we have seen in (43), the last two are already theorems of NL, whereas “Argument Raising” is not. A relevant instance of it is instead a theorem in LG: the argument of a transitive verb can be raised to the higher order category of quantifiers, as we show in (81). See (Bastenhof 2007) for further details.

As emphasized in (Szabolcsi 2008), an important task for researchers working on the syntax-semantic interface is “to determine whether the Reinhart’s hypothesis is correct and if yes, exactly what kind of syntactic representation and notion of domain bears it out”. In the logic perspective, the question reads as saying which are the logical operators that capture this connection between the two components of quantifiers, while allowing to scope over the top part of a tree and still remain in the original position as a *dp*. NL does not provide us with the required tools. It is the minimum logic to model local dependency and the assembly of direct scope constructors, but it is too weak to account for long-distance dependency, cross-dependencies and inverse scope effects. Hence, the need of a more expressive logic that goes behind these limitations. Below we show how LG properly captures the QP unbounded scope behaviour exhibited in the examples (1) and (2) above. We believe that the different scope distribution of QPs ((4) in (73)) can be accounted for by adding unary modalities to LG following (Bernardi 2002; Bernardi & Szabolcsi 2008). Following (Barker & Shan 2006), we conjecture that the cases of bounded scope ((3) in (73)) are properly handled by delimited continuations which belong to the framework we have presented here. For reason of space we cannot go into the details of continuation semantics; in Section 5 we sketch its connection with LG and refer the reader to (Bernardi and Moortgat to appear, Bastenhof 2009a).

The “parsing as deduction” view on QP Let us use an example to guide our analysis of QPs. Take the sentence *Alex thinks everyone left*. The proof of its grammaticality corresponds to the following sequent where we leave the QP category undefined.

$$(76) \quad \underbrace{dp}_{alex} \otimes (\underbrace{(dp \backslash s)/s}_{[thinks]} \otimes (\underbrace{QP}_{[everyone]} \otimes \underbrace{dp \backslash s}_{[left]})) \vdash s$$

Our task, now, is to understand which is the category to be assigned to QP. Recall the conclusion we reached while observing the quantifiers’ semantic constituent tree ((56) in Section 2.4): a quantifier consists of a syntactic and a semantic component (the *dp* and a function from *s* to *s*), such that (a) the *dp* should stay in the embedded position where it works as an argument of a verb, and (b) the scope component should be free to travel through the tree so to reach its top level, without modifying the constituent structure. In other words, we start from a structure containing a QP, that is $\Gamma [QP]$, which is proved to be of category *s*, if the simpler structure $\Gamma [dp]$ is of category *s* and the semantic components of the QP reaches the top of the structure. In the following, we will explain what it means for the QP to “reach the top” of the structure and how it does it. This behaviour is captured by the co-abstraction rule which happens within an \otimes -structure Γ containing the QP that, for the moment, we represent as $B \odot dp$ – where *B* is the category to be assigned to the QP semantic component and which still needs to be specified.

$$(77) \quad \frac{\Gamma[dp] \vdash B \oplus s}{\Gamma[B \otimes dp] \vdash s} (\otimes L)$$

In our example, $\Gamma[dp]$ stands for $dp \otimes ((dp \setminus s)/s \otimes (dp \otimes dp \setminus s))$. By application of $(\setminus L)$ and $(/L)$ the sequent is reduced to axiom links and to the simpler sequent below.

$$(78) \quad s \vdash B \oplus s$$

We want to know what could be the proper category to replace B . By dual residuation (DRES in Figure 4), we obtain that the semantic component of a QP is the co-functor category $(s \otimes s)$. We repeat in (79) the if-part of DRES we use below:

$$(79) \quad (\text{DRES}) \quad C \vdash B \oplus A \text{ iff } C \otimes A \vdash B$$

$$(80) \quad s \vdash B \oplus s \text{ iff } (s \otimes s) \vdash B$$

Hence, the whole category of a quantifier should be $(s \otimes s) \otimes dp$, where the value- s and argument- s are the first and second one, respectively, viz. $(s_{arg} \otimes s_{val}) \otimes dp$.⁹ Before going to look at some linguistic examples, it's worth underline the following two facts.

First of all, as the reader can check by herself, $(s \otimes s) \otimes dp \vdash s/(dp \setminus s)$: every expression that belongs to the non-local scope category, belongs to the local category too; whereas $s/(dp \setminus s) \not\vdash (s \otimes s) \otimes dp$. Moreover, whereas the ‘‘Argument Raising’’ is not a theorem of LG (see (43) in Section 2), the argument of a transitive verb can be raised to the QP category as shown below. In (Bastenhof 2007), this theorem is used to handle the interaction of QPs with intentional verbs.

$$(81) \quad \frac{\begin{array}{c} \vdots \\ dp \otimes ((dp \setminus s)/dp \otimes dp) \vdash (s \otimes s) \oplus s \end{array}}{dp \otimes ((dp \setminus s)/dp \otimes ((s \otimes s) \otimes dp)) \vdash s} (\otimes L)$$

$$\frac{dp \otimes ((dp \setminus s)/dp \otimes (((s \otimes s) \otimes dp))) \vdash s}{(dp \setminus s)/dp \otimes ((s \otimes s) \otimes dp) \vdash dp \setminus s} (\setminus R)$$

$$\frac{(dp \setminus s)/dp \otimes ((s \otimes s) \otimes dp) \vdash dp \setminus s}{(dp \setminus s)/dp \vdash (dp \setminus s)/(((s \otimes s) \otimes dp))} (/R)$$

Looking at the categories obtained in terms of the two merging relations, saying that $QP \in (s \otimes s) \otimes dp$ means that there are two parts y and x such that $\text{Merge}^c(y, QP, x)$, $y \notin (s \otimes s)$ and $x \in dp$. In other words, we obtain the syntactic component x by extracting the semantic component z from the QP. The syntactic component $x \in dp$ will stay in the \otimes -structure where it can merge with other expressions, whereas the semantic component will go to the \oplus -structure.

9. The s_{val} and s_{arg} of s_{val}/s_{arg} are in a positive and negative polarity position, respectively $(+/-, -/+)$ and similarly for their dual, $- \otimes +$, and $+ \otimes -$. Compare $s_1/(dp \setminus s_2)$ with $((s_2 \otimes s_1) \otimes dp)$, the polarity of the atomic formulas is the same, and s_1 is the s_{val} and s_2 is the s_{arg} . Hence, $(s_{arg} \otimes s_{val})$.


Let's check how this category behaves by looking at direct and inverse scope caused by a QP in subject (82) and object position (83), respectively. As the reader can check by herself, the application of the ($\backslash L$) (and ($/L$)) in (82) (resp. (83)) builds a derivation that does not end with axiom links –hence it fails to prove that the sequent holds. The only other possible first step is the application of ($\odot L$).

$$\begin{array}{c}
 (82) \quad \frac{\frac{\frac{s_3 \vdash s_1 \quad s_2 \vdash s_4}{s_3 \vdash (s_1 \odot s_2) \oplus s_4} (\odot R)}{dp_1 \otimes dp_2 \vdash (s_1 \odot s_2) \oplus s_4} (/L)}{dp_1 \otimes (dp_2 \backslash s_3) \vdash (s_1 \odot s_2) \oplus s_4} (\backslash L) \\
 \frac{\frac{\frac{(s_1 \odot s_2) \odot dp_1}{\text{everyone}} \otimes \frac{dp_2 \backslash s_3}{\text{left}} \vdash s_4}{((s_1 \odot s_2) \odot dp_1) \otimes (dp_2 \backslash s_3) \vdash s_4}}{((s_1 \odot s_2) \odot dp_1) \otimes (dp_2 \backslash s_3) \vdash s_4} (\odot L)
 \end{array}$$


$$\begin{array}{c}
 (83) \quad \frac{\frac{s_1 \vdash s_2 \quad s_3 \vdash s_4}{s_1 \vdash (s_2 \odot s_3) \oplus s_4} (\odot R)}{\vdots} \\
 \frac{dp_1 \otimes ((dp_2 \backslash s_1) / dp_3 \otimes dp_4) \vdash (s_2 \odot s_3) \oplus s_4}{dp_1 \otimes ((dp_2 \backslash s_1) / dp_3 \otimes ((s_2 \odot s_3) \odot dp_4)) \vdash s_4} (\odot L) \\
 \frac{[Alex] \quad [knows] \quad [everyone]}{dp_1 \otimes ((dp_2 \backslash s_1) / dp_3 \otimes ((s_2 \odot s_3) \odot dp_4)) \vdash s_4}
 \end{array}$$

With the first inference step, the co-abstraction ($\odot L$), the scope-component of the QP ($s_{arg} \odot s_{val}$) is sent to the succedent of the \vdash leaving on the antecedent a standard \otimes -structure, the syntactic constituent structure of the sentence, containing a dp in the place of the QP. By function application steps, this structure is reduced to the category s carried by its head, the verb, (s_3 in (82) and s_1 in (83)) – in (85), we will represent it with s_{vhd} . The last steps check the scope: the co-function application ($\odot R$) is applied and s_{vhd} reaches s_{arg} ($s_3 \vdash s_1$ in (82), and $s_1 \vdash s_2$ in (83)) and s_{val} reaches s_{goal} ($s_2 \vdash s_4$ in (82) and $s_3 \vdash s_4$ in (83)).

As introduced in (64) of Section 3, we use a tree for the left side of the sequent and a mirror upside down tree for its right side; we obtain the division of labor below, where the upside down tree represents the semantic constituent structure (which plays the role of May's LF) and the usual tree the syntactic structure. The first inference step ($\odot L$) in the two sequents corresponds to the schematic trees in (84) below.




To ease the comparison with the sequents above, we can represent the tree on the right in (84) as following:



where the s_{vhd} stands for the s sub-formula of the category representing the head of the predicate argument structure, viz. the s value of the main verb.

The tree, containing the dp component of QP, is reduced to s_{vhd} which reaches (\vdash) the s_{arg} of the QP's scope component and the s_{val} reaches the s_{goal}

It is interesting to compare (85) with the QR schema (75); we repeat it in (86) with the annotation of the S-nodes to help grasping the similarity between the two solutions: LG captures the mathematical principles behind QR.



The category assigned to QP built with dual-implications properly account both for direct and inverse scope. The examples above show that two of the problems encountered by the higher order categories consisting of the (Lambek) implications, viz. (i) multiple assignments and (ii) local scope, are solved. We now move to look at scope ambiguity and, in particular, at non-local scope.

The “parsing as deduction” view on Scope Ambiguity The simplest case with scope ambiguity is the one with QPs both in subject and object positions [Someone [knows everyone]]. The sequent to prove is:

$$(87) \quad \underbrace{((s_1 \otimes s_2) \otimes dp_1)}_{\text{[someone]}} \otimes \underbrace{((dp_2 \backslash s_3) / dp_3)}_{\text{[knows]}} \otimes \underbrace{((s_4 \otimes s_5) \otimes dp_4)}_{\text{[everyone]}} \vdash s_6$$

There are three main functional connectives (two \otimes and one $/$) in the antecedent of the sequent that could be activated, hence three inference rules could be applied. As the reader can check by herself, the application of the $(/L)$ rule builds a derivation that does not end with axiom links –hence it fails to prove that the sequent holds. The other two possibilities are $(\otimes L)$ applied to the \otimes of the subject or of the object. These two choices bring to the two different derivations schematized below.

$$\begin{array}{c}
(88) \quad \begin{array}{c} s_5 \vdash s_1 \quad s_2 \vdash s_6 \quad s_3 \vdash s_4 \\ \vdots \\ s_3 \vdash (s_4 \odot s_5) \oplus ((s_1 \odot s_2) \oplus s_6) \\ \vdots \\ \frac{dp_1 \otimes ((dp_2 \setminus s_3)/dp_3 \otimes dp_4) \vdash (s_4 \odot s_5) \oplus ((s_1 \odot s_2) \oplus s_6)}{dp_1 \otimes ((dp_2 \setminus s_3)/dp_3 \otimes ((s_4 \odot s_5) \odot dp_4)) \vdash (s_1 \odot s_2) \oplus s_6} (\odot L) \\ \frac{dp_1 \otimes ((dp_2 \setminus s_3)/dp_3 \otimes ((s_4 \odot s_5) \odot dp_4)) \vdash (s_1 \odot s_2) \oplus s_6}{((s_1 \odot s_2) \odot dp_1) \otimes ((dp_2 \setminus s_3)/dp_3 \otimes ((s_4 \odot s_5) \odot dp_4)) \vdash s_6} (\odot L) \\ \underbrace{\hspace{10em}}_{\text{[someone]}} \quad \underbrace{\hspace{10em}}_{\text{[knows]}} \quad \underbrace{\hspace{10em}}_{\text{[everyone]}} \end{array} \\
\text{[SOME > EVERY]}
\end{array}$$

$$\begin{array}{c}
\begin{array}{c} s_5 \vdash s_6 \quad s_2 \vdash s_4 \quad s_3 \vdash s_1 \\ \vdots \\ s_3 \vdash (s_1 \odot s_2) \oplus ((s_4 \odot s_5) \oplus s_6) \\ \vdots \\ \frac{dp_1 \otimes ((dp_2 \setminus s_3)/dp_3 \otimes dp_4) \vdash (s_1 \odot s_2) \oplus ((s_4 \odot s_5) \oplus s_6)}{(s_1 \odot s_2) \odot dp_1 \otimes ((dp_2 \setminus s_3)/dp_3 \otimes ((s_4 \odot s_5) \odot dp_4)) \vdash (s_4 \odot s_5) \oplus s_6} (\odot L) \\ \frac{(s_1 \odot s_2) \odot dp_1 \otimes ((dp_2 \setminus s_3)/dp_3 \otimes ((s_4 \odot s_5) \odot dp_4)) \vdash (s_4 \odot s_5) \oplus s_6}{(s_1 \odot s_2) \odot dp_1 \otimes ((dp_2 \setminus s_3)/dp_3 \otimes ((s_4 \odot s_5) \odot dp_4)) \vdash s_6} (\odot L) \\ \underbrace{\hspace{10em}}_{\text{[someone]}} \quad \underbrace{\hspace{10em}}_{\text{[knows]}} \quad \underbrace{\hspace{10em}}_{\text{[everyone]}} \end{array} \\
\text{[EVERY > SOME]}
\end{array}$$

Again, the structure $dp_1 \otimes ((dp_2 \setminus s_3)/dp_3 \otimes dp_4)$ reduces to s_3 by function application steps as explained in Section 2, the reader is invited to check the steps by herself. Instead, we spell out the second part that brings to the relevant axioms by means of co-function applications:¹⁰

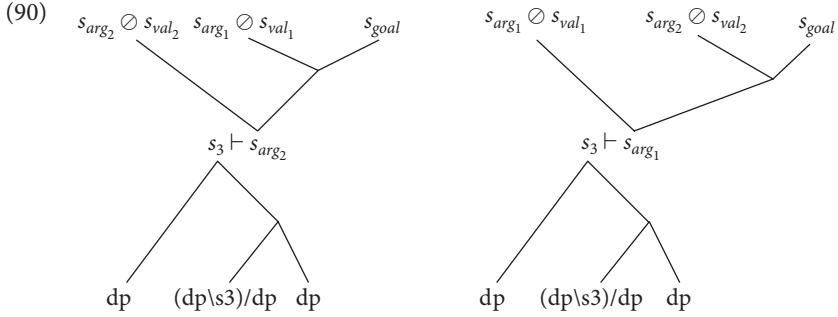
$$\begin{array}{cc}
(89) & \begin{array}{c} \text{[SOME > EVERY]} \\ \frac{s_5 \vdash s_1 \quad s_2 \vdash s_6}{s_5 \vdash (s_1 \odot s_2) \oplus s_6} (\odot L) \\ \frac{s_5 \vdash (s_1 \odot s_2) \oplus s_6 \quad s_3 \vdash s_4}{s_3 \vdash (s_4 \odot s_5) \oplus ((s_1 \odot s_2) \oplus s_6)} (\odot L) \end{array} & \begin{array}{c} \text{[EVERY > SOME]} \\ \frac{s_5 \vdash s_6 \quad s_2 \vdash s_4}{s_2 \vdash (s_4 \odot s_5) \oplus s_6} (\odot L) \\ \frac{s_2 \vdash (s_4 \odot s_5) \oplus s_6 \quad s_3 \vdash s_1}{s_3 \vdash (s_1 \odot s_2) \oplus ((s_4 \odot s_5) \oplus s_6)} (\odot L) \end{array}
\end{array}$$

Recall the observation made above regarding the s -value and s -argument of the QP, viz. $(s_{arg} \odot s_{val}) \odot dp$. As shown by the axiom links in the first derivation the s -value of *someone* matches the s -goal formula ($s_2 \vdash s_6$), the s -value of *everyone* matches the s -argument of *someone*, and the s of the transitive verb matches the s -argument of *everyone*. These are the axiom links of the subject wide scope reading [SOME > EVERY]. The axioms linking s -categories of the other derivation are different: they give the object wide scope reading [EVERY > SOME]. Observe that the earliest a quantifier jumps out of the \otimes -structure, the widest scope it has (recall we are looking at scope through a mirror! –through the dual residuation calculus). Notice, that the axioms linking dp -categories are the same in the two derivations, i.e. $dp_1 \vdash dp_2$ and $dp_4 \vdash dp_3$; in both derivations,

10. Notice that in the dual calculus too, we could follow either a top-down or a bottom-up approach to derive the sequents. For instance, in (88) we could have applied first the internal function to match the s_6 formula first.

the variable taken as object by the verb is bound by the QP in object position and the variable taken as subject of the verb is bound by the QP in subject position.

In a tree notation the two derivations above are schematized as in (90). Again, to help thinking of *s*-value and *s*-argument, here we mark the *s*-categories with this information instead of numbers. $s_{arg_2} \otimes s_{val_2}$ is the semantic component of the Quantifier in object position, *everyone*, i.e. $s_4 \otimes s_5$; whereas $s_{arg_1} \otimes s_{val_1}$ is the semantic component of the quantifier in subject position, *someone*, i.e. $s_1 \otimes s_2$.



The two readings are read out of the upside down trees: the tree on the left gives the subject wide scope reading, and the tree on the right gives the object wide scope reading. The syntactic constituents are read out of the usual tree that is the same in the two cases – it is the sequent that reduces to s_3 in the two derivations above. The cases with QP interacting with other scope operators will work out in a similar way. The reader is invited to try, for instance, to build the derivations for the “John doesn’t read a book” (viz. prove that $dp \otimes ((dp\s)/s) \otimes ((dp\s)/dp \otimes ((s \otimes s) \otimes dp)) \vdash s$).


The case of unbounded scope works in the same way: given the structure corresponding to [Alex [thinks [everyone left]]], there are two main connectives / and \otimes , hence there are two possible inference rules to be applied as a first step: Either we activate first *everyone* (91) or *thinks* (94).

$$\begin{array}{c}
 (91) \quad \frac{s_4 \vdash s_6 \quad s_1 \vdash s_3}{s_1 \vdash (s_3 \otimes s_4) \oplus s_6} (\otimes R) \\
 \vdots \\
 \frac{(dp \otimes ((dp\s_1)/s_2 \otimes (dp \otimes dp\s_5))) \vdash (s_3 \otimes s_4) \oplus s_6}{[EVERY > THINK] \quad \frac{dp_1 \otimes ((dp_2\s_1)/s_2 \otimes ((s_3 \otimes s_4) \otimes dp_3) \otimes dp\s_5) \vdash s_6}{\begin{array}{cccc} \text{Alex} & \text{[thinks]} & \text{[everyone]} & \text{left]} \end{array}}} (\otimes L)
 \end{array}$$

Again, the full structure on the left of the \vdash , by means of function applications, reduces to s_1 (i.e., the value carried by the head of the structure, *thinks*); hence, as we did above, we can abbreviate the steps of the derivation and focus on the branch of the derivation containing the dual residuated operators and that account for the scope structure. The scope component of the quantifier has reached the wide position


of the whole structure, and takes scope on its semantic domain on which it is applied by means of co-function application ($\odot R$). The value- s (s_4) of the operator (*everyone*) having wide scope matches the s -goal formula (s_6), whereas its argument- s (s_3) is obtained from the value- s (s_1) of the narrow operator (*thinks*). We can represent the sequents' steps with the following tree. Now that the reader might be familiar with the procedure, we take the liberty of using words instead of the category and give only the schema of the trees.

(92)



In the second reading of the sentence, the syntactic and semantic domain of *thinks* coincides: it takes scope over the structure it c-commands/dominates, hence also on the quantifier. Proof theoretically this means that its s -value (s_1) is linked to the goal formula (s_6), and its s -argument (s_2) is reached by the quantifier s -argument (s_4). The application of the function *thinks* to the expression *everyone left* reduces the problem of verifying that the full structure is of category s_6 to two simpler problems: verifying that [everyone left] is s_2 and the subject dp_1 composed with the verb phrase category $dp_1 \setminus s_1$ is of category s_6 . The first tree is checked as we have seen in (82), the second simply by function application.

(93)




The derivation corresponding to this second reading is built by focusing first on *thinks* and then on *everyone*; as we have seen in the case of multiple quantifiers, the operator which is activated earlier receives wider scope.

(94)

$$\begin{array}{l}
 \frac{dp_3 \vdash dp_4 \quad s_5 \vdash s_3 \quad s_4 \vdash s_2}{dp_3 \vdash dp_4 \quad s_5 \vdash (s_3 \odot s_4) \oplus s_2} (\odot L) \\
 \frac{dp_3 \otimes dp_4 \setminus s_5 \vdash (s_3 \odot s_4) \oplus s_2}{((s_3 \odot s_4) \odot dp_3) \otimes dp_4 \setminus s_5 \vdash s_2} (IL) \\
 \frac{dp_1 \otimes ((dp_2 \setminus s_1) / s_2) \otimes (((s_3 \odot s_4) \odot dp_3) \otimes dp_4 \setminus s_5) \vdash s_6}{\text{Alex} \quad [\text{thinks} \quad [\text{everyone} \quad \text{left}]]} (\odot L) \quad \frac{dp_1 \vdash dp_2 \quad s_1 \vdash s_6}{dp_1 \otimes dp_2 \setminus s_1 \vdash s_6} (\setminus L)
 \end{array}$$


Summing up, with (82) and (83) we have shown that the category assigned to the QP using the extended language (i) avoids multiple assignments, allows for (ii) local scope; this last example shows that it also handles (iii) non-local scope. Moreover, the examples with multiple scope operators show that syntactic structures containing scope operators can be built by activating them in different orders and these different steps produce different readings of the same structure. Finally, in LG we exploit its two relations on trees to build the constituent tree (c-tree) and the scope commanded tree (s-tree), when the latter may diverge from the former: The c-tree is built out of \otimes whereas the s-tree is built out of \oplus , as illustrated by the tree on the right in (95). A scope operator that acts only locally, like *thinks*, scopes on the structure it c-commands, a \otimes -structure: by means of function applications its s-value reaches the goal-s formula, as illustrated by the tree on the left in (95):

(95)



If we go back to the preliminary analysis done in Section 2.4, when we introduced the concept of a semantic and syntactic component, we see that the tree we proposed there can be made precise. We can visualize the combination of the tree and upside down tree given in (90) as below to ease the comparison. Let's take the reading with $OP_1 > OP_2$, the tree looks like the following

(96)



We can now go back to the goal of this paper, viz. use the “parsing as deduction” approach to shed new lights on the question about Reinhart’s hypothesis, repeated below in (97). The question we are seeking to answer is: “is this hypothesis correct?

And if yes, exactly what kind of syntactic representation and notion of domain bears it out” (Szabolcsi 2008).

- (97) “The scope of a linguistic operator coincides with its domain in some syntactic representation that the operator is part of.”

Our answer to this question is that, yes, Reinhart’s hypothesis is correct, and that scope operators require to distinguish their syntactic domain from their semantic domain which are governed by dual relations.

Adding the derivational perspective worked out in this paper, we could rephrase (97) in the following way

- (98) “The scope of a linguistic operator is *read off* of the (inference) *steps* that put together the expression of which the operator is part of.”

5. Related works and some further intuition

Two points of our work that is most interesting to see in connection with related works are the idea of looking at quantifier as consisting of a syntactic and a semantic component and the interpretation of natural language via Continuation Semantics.

5.1 QP syntactic category

The idea of looking at quantifiers as consisting of a syntactic and a scope component is not new. Examples of its use, among others, are the work by Barker and Shan (Barker & Shan 2008), and the one by Kallmeyer and Joshi (Kallmeyer & Joshi 2003). In the former the authors give a nice and clear account of donkey anaphora relating it to quantifier binding. The grammar they use, based on continuation semantics, had been previously introduced in (Barker & Shan 2006), but in (Barker & Shan 2008) is presented with a simplified “tower” notation that highlights the different role played by the *dp*-component of the QP and its scope-part. A quantifier receives the category below, that, the author says, “can be read counterclockwise, starting below the horizontal line” –here and in the following brief overview, we indicate the value-*s* and the argument-*s* for ease of comparison with our proposal. For a comparison between our and Barker and Shan’s proposal see (Bastenhof 2009a).

- (99) $\frac{s_{val}|s_{arg}}{dp}$ means $\frac{\dots \text{ to form an S. | and takes scope at an S}}{\text{The expression functions in local syntax as a DP}}$

Similarly, in (Kallmeyer & Joshi 2003), the authors extend Tree Adjoining Grammar with multi-component trees in order to deal with QP scope ambiguity. In this framework, lexical entries receive a tree; a QP receives a tree with two components: “one component

corresponds to the contribution of [the quantifier] to the predicate-argument structure [...], and [the other] component contributes to the scope structure.”

Finally, Moortgat (Moortgat 1996) proposes a three-place binding type constructor q which captures the behavior of in-situ binders. The behaviour of this operator was represented by the inference step below. An expression of category $q(A, B, C)$ occupies the position of an expression of category A within a structural context of category B ; using the q connective turns the domain of category B into one of type C . In particular, a QP receives the category $q(dp, s_{arg}, s_{val})$.

$$(100) \quad \frac{\Delta[A] \vdash B \quad \Gamma[C] \vdash D}{\Gamma[\Delta[q(A, B, C)]] \vdash D}$$

This inference step properly characterizes natural language scope operators. However, this operator is not part of the logical language, and it was not known which are the logical operators that capture its behaviour. Our work is an answer to this question. Notice, the limitations encountered by the q -operators discussed in (Carpenter 1998), e.g. its relation with coordination, are overcome by the proposal presented in this paper.

The idea of looking at QPs has an effect of splitting a given expressions into parts is also behind the work in (Morrill et al. 2007; Morrill & Fadda 2008). In these papers the starting point is the Associative Lambek calculus (Lambek 1961) which is extended with discontinuous product and their residuals, infixation (\downarrow) and extraction (\uparrow), and separators to mark the holes generated by the extraction and where other expression can be substitute in. In this framework, called Discontinuous Lambek Calculus, a QP receives the category $(s_{arg} \uparrow dp) \downarrow s_{val}$: the infixation substitutes the QP in the hole generated by the extraction of a dp from the predicate-argument structure; the hole is marked by a separator. It is interesting to notice that this extension has been shown to overcome the limitations mentioned in Section 1, being able to account both for quantifier scope construals and non-peripheral extractions and cross-serial dependencies. See (Bastenhof 2009b) and (Moot 2007) for a discussion of these phenomena in LG.

We would like to draw attention on an important difference between the solutions mentioned so far. In Morrill and Fadda and Barker and Shan’s categories the dp and s_{arg} form a sub-formula together in the QP category, whereas both in Kallmeyer and Joshi and in our’s category the dp category is taken apart. We believe that is a point of interest that deserves further in depth analysis.

5.2 Continuation semantics


Barker and Shan, and Philippe de Groote have shown in several papers (see for instance (Barker 2002; Shan 2005; Barker & Shan 2006; Barker & Shan 2008; De Groote 2001;

De Groote 2006)) that continuation semantics is what is needed to interpret quantifier phrases as well as other related phenomena, like donkey anaphora, wh-extraction etc. Important keywords of this semantics are “values”, “continuation” and “computations”. Below we try to give an intuitive idea of their difference and we briefly discuss how continuation semantics properly interprets the LG derivations discussed so far.

Values, Continuation and Computation Work done within Curry Howard Correspondence to Classical Logic has brought to the attention the duality distinction holding between terms and contexts. We believe that this “duality” view can help further shaping the syntax-semantics interface and the connection between natural and formal languages. This duality is clearly explained in (Sacerdoti Coen 2006):

- A context (aka, *continuation*) is an expression with one placeholder for a “missing” term (aka, *value*). The placeholder can be filled with a term (a value) to obtain a placeholder-free expression. ($K_A V_A = R$)
- Dually, a term (aka, *computation*) can be seen as an expression with exactly one placeholder for a “missing context” that is “all around” the term. The placeholder can be filled with a context to obtain a placeholder-free expression. ($C_A K_A = R$)

We will abbreviate *values*, *continuation* and *computation* of type A as V_A , K_A and C_A , respectively. We can think of “a placeholder-free expression” as an “observable object” and assign to it the type R (the category of an unfocused entailment when the communication between the two merge-relations is established, i.e. the conclusion of a Cut-rule (See example in (67)). Both the application of a continuation to a value ($K_A V_A$) and of a computation to a continuation ($C_A K_A$) result into a “placeholder-free expression” (R). Since we are interested in composing linguistic structures (tree-structures), the above applications can be visualized as below. A value is a tree, a continuation and a computation are a tree with a hole on the bottom and on the top part, respectively.



The picture highlights the new view on compositionality, and its connection with the two merging relations discussed so far might be clear at this point: the \otimes composes values whereas the \oplus composes continuations. The $\bar{\lambda}\mu\tilde{\mu}$ -calculus in Curry Howard Correspondence with LG has the μ and $\tilde{\mu}$ operator to bring the focus on elements of the \oplus and \otimes structures, respectively.

6. Conclusion

Our goal was to define the mathematical principles of natural language structures. We have claimed that the fundamental principles needed are the (pure) residuation principle, its dual and the distributivity principles which establish the interaction among them. We have shown that, proof-theoretically, this means having structures both on the antecedent and the succedent of the entailment relation of a sequent. We have also briefly discussed the connection between this extended calculus and Continuation Semantics, that, in several papers by different authors, has been shown to be a proper framework for the interpretation of natural language structures containing inverse scope operators and the alike.

References

- Areces, C. & Bernardi, R. 2004. Analyzing the core of categorial grammar. *Journal of Logic, Language and Information*, 13–2 (2004), 121–137.
- Barker, C. 2002. Continuations and the nature of quantification. *Natural Language Semantics*, 10 (2002), 211–242.
- Barker, C. & Shan, C. 2006. Explaining crossover and superiority as left-to-right evaluation. *Linguistics and Philosophy*, 29–1 (2006), 91–134.
- Barker, C. & Shan, C. 2008. Donkey anaphora is in-scope binding. *Semantics and Pragmatics*, 1 (2008), 1–46.
- Bastenhof, A. 2007. Quantifier scope and coordination in flexible and continuation-based approaches to natural language semantics. B.Phil thesis, Utrecht University.
- Bastenhof, A. 2009a. Continuation in Natural Language Syntax and Semantics. M.Phil thesis, Utrecht University.
- Bastenhof, A. 2009b. Extraction in the Lambek-Grishin calculus. In *Proceedings of the 14th Student Session of the European Summer School for Logic, Language, and Information*, Thomas Icard, ed., p. 106–126, Bordeaux, ESLLI.
- Benthem, J. van. 1987. Categorial Grammar and Lambda Calculus. In *Mathematical Logic and its Applications*, D. Skordev, ed., p. 39–60. New York, Plenum.
- Benthem, J. van. 1988. The Lambek Calculus. In *Categorial Grammars and Natural Language Structures*, R. Oehrle, E. Bach & D. Wheeler, ed., p. 35–68. Dordrecht, Reidel.
- Bernardi, R. 2002. *Reasoning with Polarity in Categorial Type Logic*. Utrecht, Utrecht Institute of Linguistics OTS.
- Bernardi, R. & Moortgat, M. to appear. Continuation semantics for the Lambek-Grishin calculus. *Information and Computation*, (to appear).
- Bernardi, R. & Moortgat, M. 2007. Continuation Semantics for Symmetric Categorial Grammar. In *Proceedings of the 14th Workshop on Logic, Language, Information and Computation (WOLLIC'07)*, D. Leivant & R. de Queiroz, ed., LNCS 4576, p. 53–71, Berlin/Heidelberg, Springer.
- Bernardi, R. & Szabolcsi, A. 2008. Optionality, scope, and licensing: An application of partially ordered categories. *Journal of Logic, Language and Information*, 17–3 (2008).

- Carpenter, B. 1998. *Type-Logical Semantics*. Cambridge, MA, MIT Press.
- Chomsky, N. 1976. Conditions on Rules of Grammar. *Linguistic Analysis*, 2 (1976), 303–351.
- Goré, R. 1997. Substructural Logics on Display. *Logic Journal of IGPL*, 6–3 (1997), 451–504.
- Grishin, V.N. 1983. *Studies in Nonclassical Logics and Formal Systems*, chapter On a generalization of the Ajdukiewicz-Lambek system, p. 315–334. Moscow, Nauka. English translation in Abrusci and Casadio (eds.) *Proceedings 5th Roma Workshop*, Bulzoni Editore, Roma, 2002.
- Groote, P. de, 2001. Type raising, continuations, and classical logic. In *Proceedings of the Thirteenth Amsterdam Colloquium*, R. van Rooy & M. Stokhof, ed., p. 97–101, Amsterdam, ILLC, University of Amsterdam.
- Groote, P. de, 2006. Towards a Montagovian account of dynamics. In *Proceedings of Semantics and Linguistic Theory XVI*, CLC Publications.
- Hendriks, H. 1993. *Studied Flexibility. Categories and Types in Syntax and Semantics*. Amsterdam, ILLC, University of Amsterdam.
- Howard, W.A. 1980. The formulae-as-types notion of construction. In *To H.B. Curry: Essays on Combinatory Logic, Lambda Calculus and Formalism*, p. 480–490. London, Academic Press.
- Kallmeyer, L. & Joshi, A. 2003. Factoring Predicate Argument and Scope Semantics: Underspecified Semantics with LTAG. *Research on Language and Computation*, 1-1-2 (2003), 3–58.
- Kurtonina, N. & Moortgat, M. 2007. Relational semantics for the Lambek-Grishin calculus. In *10th Conference on Mathematics of Language*, M. Kracht, G. Penn & E. Stabler, ed., Los Angeles, UCLA Working Papers in Linguistics.
- Lambek, J. 1958. The Mathematics of Sentence Structure. *American Mathematical Monthly*, 65 (1958), 154–170.
- Lambek, J. 1961. On the calculus of syntactic types. In *Structure of Languages and its Mathematical Aspects*, R. Jakobson, ed., p. 166–178. American Mathematical Society.
- Lambek, J. 1993. From categorial to bilinear logic. In *Substructural Logics*, K. Došen & P. Schröder-Heister, ed., p. 207–237. Oxford, Oxford University Press.
- May, R.C. 1977. *The Grammar of Quantification*. Cambridge, MA, Massachusetts Institute of Technology.
- Montague, R. 1974. *Formal Philosophy: Selected Papers of Richard Montague*, R. Thomason (ed.). New Haven, Yale University Press.
- Moortgat, M. 2009. Symmetric Categorical Grammar. *Journal of Philosophical Logic*, 681–710.
- Moortgat, M. 1988. *Categorical Investigations. Logical and Linguistic Aspects of the Lambek calculus*. Dordrecht, Foris.
- Moortgat, M. 1996. Generalized quantifiers and discontinuous type constructors. In *Discontinuous Constituency*, H. Bunt & A. van Horck, ed., p. 181–207. Berlin/New York, De Gruyter.
- Moortgat, M. 2007. Symmetries in natural language syntax and semantics: the Lambek-Grishin calculus. In *Proceedings 14th Workshop on Logic, Language, Information and Computation (WoLLIC'07)*, D. Leivant & R. de Queiros, ed., LNCS 4576, p. 264–284, Berlin/Heidelberg, Springer.
- Moot, R. 2007. Proof nets for display logic. CoRR abs/0711.2444.
- Morrill, G. & Fadda, M. 2008. Proof nets for basic discontinuous Lambek calculus. *Journal of Logic and Computation*, 18–2 (2008), 239–256.
- Morrill, G. Fadda, M. & Valentin, O. 2007. Nondeterministic Discontinuous Lambek Calculus. In *Proceedings of the Seventh International Workshop on Computational Semantics (IWCS7)*, Tilburg.

- Reinhart, T. 1979. Syntactic domains for semantic rules. *Formal Semantics and Pragmatics for Natural Languages*, (1979), 107–130.
- Reinhart, T. 1997. Quantifier Scope: How Labor is Divided Between QR and Choice Functions. *Linguistic and Philosophy*, 20–4 (1997), 335–397.
- Ruys, E. & Winter, Y. to appear. *Handbook of Philosophical Logic*, chapter Quantifier Scope in Formal Linguistics.
- Sacerdoti Coen, C. 2006. Explanation in Natural Language of $\bar{\lambda}\mu\bar{\mu}$ terms. In *Fourth International Conference on Mathematical Knowledge Management, LNAI*, volume 3863, Berlin/Heidelberg, Springer.
- Schröder, E. 1980. *Vorlesungen über die Algebra der Logik*, volume 1. Leipzig, Teubner.
- Shan, C. 2005. *Linguistic Side Effects*. Harvard University.
- Steedman, M. 2000. *The Syntactic Process*. Cambridge, MA, MIT Press.
- A. Szabolcsi, ed., 1997. *Ways of Scope Taking*. Berlin/Heidelberg, Springer.
- Szabolcsi, A. 2008. Scope and Binding. In *Semantics: An International Handbook of Natural Language Meaning*, C. Maienborn, K. von Stechow & P. Portner, ed., HSK series. Berlin, Mouton de Gruyter.
- Vermaat, W. 2005. *The logic of variation. A cross-linguistic account of wh-question formation*. Utrecht, UiL OTS, Utrecht.

Phonetic and phonological approaches to early word recognition

Empirical findings, methodological issues, and theoretical implications

Paola Escudero & Titia Benders
University of Amsterdam

This paper reviews two methodological paradigms and three theoretical proposals put forward to investigate infants' sensitivity to phonetic detail in early word recognition. The results of studies using the first methodological paradigm suggest that infants are unable to use phonetic detail when accessing the representations of newly learned words. In contrast, infants tested using the second methodological paradigm generally succeed at identifying small differences between words. The three theoretical proposals differ in the extent to which they can explain the rate of success in infant word recognition. Methodological adaptations and new research directions to bridge the gap between methodologies and theories are presented.

1. Introduction

It is well known that young infants are able to distinguish between almost all sound contrasts on which they have been tested (cf. Aslin, Jusczyk & Pisoni 1998, for a review), and have even been shown to surpass the discriminating abilities of adult listeners (Werker, Gilbert, Humphrey & Tees 1981). However, this ability weakens at around 10 months, if the task is to distinguish non-native consonants (Werker & Tees 1984), while it improves at around the same age, if the task is to discriminate native consonants (Kuhl, Stevens, Hayashi, Deguchi, Kiritani & Iverson 2006). With regards to vowel perception, the tuning towards the native language categories appears about four months earlier than for consonants (Kuhl, Williams, Lacerda, Stevens & Lindblom 1992; Polka & Werker 1994).

One could then ask how these early changes in phonetic discrimination impact language learning, specifically the storage and recognition of word forms. That is, is it the case that infants exploit their language specific phonetic knowledge, which is in place after the first year of life, for purposes of early word learning in the second

year of life? In this paper, we review a recent tradition of infant word recognition studies, which aims at addressing this issue using two main methodological paradigms. The two methodological paradigms can be referred to as the *switch* and the *mispronunciation* procedures.

In Section 2, we first review the early switch studies which have investigated word learning in infants as young as 14 months of age. Then, we present a phonological proposal which aims at explaining early word learning, and discuss its adequacy in accounting for the data. Two alternative phonetic-phonological approaches will be put forward to account for the results which the first approach cannot fully explain. In Section 3, we present other switch studies which show that infants can have more success in learning minimally different words, if the task is simple enough for them. Here, we show that the switch procedure may not be the most sensitive procedure to investigate how infants use phonetic detail in early word recognition.

The fourth section presents the other main paradigm used to investigate early word recognition, namely the mispronunciation procedure. There have been a number of studies which investigate how infants react to the incorrect pronunciations of familiar words and of *newly-learned* (also called *novel*) words. Although most studies within this paradigm report on infants' behaviour when presented with *mispronounced* words which vary greatly from their correct pronunciations, this section will concentrate on recent studies which aim at comparing the effects of minimal differences between correct and incorrect pronunciations. Importantly, our goal is not to be exhaustive in the presentation of all the studies conducted within each paradigm. Rather, we aim at presenting the standpoints and showing the shortcomings of the current findings and explanations.

In the last section, we outline directions for future research that are necessary for an ultimate understanding of the interrelation between speech perception, word learning and word representation. We introduce a methodological adaptation which aims at rigorously testing how infants process and encode small phonetic differences when learning minimally different words. This method allows to investigate the difference between sound discrimination and categorization and differences between age groups.

2. Research on early word recognition using the switch paradigm

The switch procedure measures infants' ability to rapidly learn the association between auditorily presented novel labels and pictures of novel objects. This procedure was first introduced by Werker, Cohen, Lloyd, Casasola and Stager (1998). In the training phase, infants are habituated to two label-object pairings. In the test phase, the crucial trial is the *switch trial*, in which object A, which was originally paired with label A, appears

now paired with label B or vice versa. Infants' looking times to the switch trials, which violate the habituated label-object associations, are compared to their looking times to the *same trials*, which maintain the habituated label-object associations. If children look longer at switch than at same trials, it means that they notice the violation of the habituated association and, consequently, that they must have learned the label-object associations. Using this procedure, Werker et al. (1998) showed that 14-month-old infants are able to learn to associate two dissimilar sounding words (i.e. /lɪf/ and /nɪm/) to two different objects in a very short period of time. Apparently, at this age, infants can learn arbitrary associations between objects and auditorily presented labels, which is an ability crucial to word learning.

Despite 14-month-olds' ability to learn word-object associations involving two dissimilar sounding words, Stager and Werker (1997) found that infants of the same age failed at detecting the switch of previously learned label-object pairings when the minimally different words /bɪ/ and /dɪ/ were presented. This lack of a switch effect was found even in a simplified version of the switch procedure in which infants are taught only one label-object association (e.g. /bɪ/-object) and the switch trial consists of the learned object combined with a new label (e.g. /dɪ/-object). This failure of 14-month-olds to detect the switch of two minimally different words has been replicated numerous times, using the same and different stop consonants (Werker, Corcoran, Fennell & Stager 2002; Pater, Stager & Werker 2004; summarized in Table 1).

Table 1. Fourteen month-old infants' results in the switch procedure: Phonological features and auditory cues signaling the consonant contrast difference for the two learned words

Contrast	Feature	Main auditory cue	Detection of Switch	Study
/b/ vs. /d/	Place	Formant transition	Failed	Stager and Werker (1997), Werker et al. (1998), Pater et al. (2004)
/d/ vs. /g/	Place	Formant transition	Failed	Pater et al. (2004)
/b/ vs. /p/	Voicing	VOT ¹	Failed	Pater et al. (2004)
/d/ vs. /p/	Place + Voicing	Formant transition + VOT ²	Failed	Pater et al. (2004)

As shown in Table 1, infants at 14 months of age fail to notice the switch when learning words contrasting in their initial consonant regardless of whether one or two

1. VOT = Voice Onset Time.
2. VOT = Voice Onset Time.

features differentiate the consonants. If the amount of dissimilarity drove success in word recognition, at this age, one would imagine that infants could show more success at the contrast /d/-/p/ than at the others, because it involves a larger difference, measured in terms of phonological features or auditory cues. In addition, if infants found particular contrasts easier than others, we would potentially see differences between contrasts involving place of articulation or formant transitions and contrasts involving voicing or voice onset time. However, previous studies show that, as a group, infants detect the switch between words that differ minimally in their consonants at 17 months of age but not before (Werker et al. 2002).

What is important to keep in mind is that infants of this age have no trouble differentiating the word forms /dɪ/ and /bɪ/ when the task is a speech discrimination task and does not involve any label-object association (Stager & Werker 1997). One of the important questions regarding early language development is what causes infants to apparently not use their native-language perceptual abilities when the task is to store and retrieve new words.

Fikkert (to appear) reviews a number of studies which were conducted to replicate previous findings and to examine early word recognition in a language other than English. The experiments tested Dutch 14-month-old infants' ability to detect the switch between the consonants /b/ and /d/ in two different vowel contexts, i.e. /bɪn/-/dɪn/ and /bɔn/-/dɔn/. The experiments were conducted using the simple version of the switch procedure, in which only one word is learned during the habituation phase. The results show that infants noticed the switch between /b/ and /d/ when they learned either /bɔn/ or /dɔn/, irrespective of whether they learned /bɔn/ and heard the novel word /dɔn/ on the switch trial or vice versa. In contrast, the infants failed to notice the switches between /bɪn/ and /dɪn/. Therefore, it seems that infants are able to notice the switch between /b/ and /d/ but only in certain contexts. These results are particularly interesting because they suggest that the English-speaking infants of previous studies could have succeeded at detecting the switch between the consonants of novel words, if such consonants had been followed by a different vowel.

Fikkert (to appear) reports the first switch studies where the vowels of words, rather than their consonants, are substituted in switch trials. Again using the simple version of the switch paradigm, this experiment tested infants' ability to notice the switch between /dɪn/ and /dɔn/ and that between /bɪn/ and /bɔn/. Here it was found that infants did not notice the switch between /dɪn/ and /dɔn/. In contrast, they noticed the switch between /bɪn/ and /bɔn/ but only when the infants learned the word /bɔn/ during the habituation phase. These results suggest not only that the context in which a contrast is presented influences the detection of the switch, but also that the direction of the switch is important. Fikkert (to appear) provides a phonological explanation for these puzzling results, which we review in the next section.

Curtin, Fennell and Escudero (2009) constitutes the only switch study which tests vowel contrasts using the more complex version of the switch procedure, in which two words are presented during the habituation phase. It is also the first study which considers more than one vowel contrast. They tested three groups of 15-month-old Canadian English infants, who were exposed to one of three word pairs, namely /dit/-/dut/, /dit/-/dɪt/, and /dɪt/-/dut/. Thus, for instance, in the /dit/-/dut/ group, all infants learned to associate the label /dit/ to object A, and the label /dut/ to object B. At test, the switch trial for half of the infants was object A paired with the label /dut/, and it was object B paired with the label /dit/, for the other half.

Figure 1 shows a representation of the looking times to the same and switch trials.³ Recall that a significant difference between infants' looking times to the two types of trials demonstrates their ability to detect the switch in the learned word-object associations and therefore their ability to recognize that the words are different. As shown in the figure, the infants who learned the pair /dit/-/dɪt/ were the only ones who detected the switch of the vowels.




Figure 1. Looking time on the same and switch trials for the contrasts /dit/-/dɪt/, /dit/-/dut/ and /dut/-/dɪt/, adapted from Curtin et al. (2009). * = $p < 0.05$

3. For each pair, the results from the switches in both directions are collapsed. Importantly, Curtin et al. did not find an effect for the direction of the switch in any of the three vowel contrasts (Curtin, personal communication).

These results show again that the detection of the switch is influenced by the specific sound contrast with which infants are presented. In the next section, we will discuss whether three proposals which aim at explaining early word recognition can account for infants' differential performance in these three vowel contrasts which are embedded in the same consonantal context, i.e. /d-vowel-t/.

2.1 Phonological explanation

One explanation for the fact that 14-month-old infants fail to notice some switches between newly learned similar sounding words, but not others, can be found within the *Featural Underspecified Lexicon* (FUL) model developed by Lahiri and Reetz (2002). This model was first applied to the representation of children's first words by Fikkert (to appear). In short, it is proposed that young children's representations are global and underspecified. In global representations, the representation of a whole word is determined by the features of one segment, the place of articulation (further: PoA) of the vowel (following Levelt 1994; and Fikkert & Levelt 2008). For instance, if children learn the word /dɔn/, its lexical representation will be *labial* because the vowel is labial, despite the fact that the onset consonant has a *coronal* feature.⁴ This global phonological representation does not prevent the infant from perceiving the PoA of the other segments. On the other hand, the underspecification of representations means that, even some vowel features, although phonetically perceived by infants, are not phonologically represented in their lexicon. Within the proposal, it is assumed that the feature *coronal*, which is found in the perceived form of the vowel in, for instance, /dɪn/, is underspecified in the lexicon. This hypothesis is based on the assumption that *coronal* is the default place of articulation. The authors mention that the strongest evidence for the default status of *coronal* within young children's representations comes from child production data, where coronal consonants substitute other consonants, at specific developmental stages (Fikkert & Levelt 2008: Footnote 10).⁵ In general, this means that there is a discontinuity between speech perception⁶ and lexical representation, i.e. not all phonetically perceived features are stored in the phonological representation (but see Section 2.2 for a different approach on perceptual and lexical representations).

4. The final consonant is not taken into account because of its very low saliency.

5. The underspecification of coronal is supported by a large number of studies (cf. Paradis & Prunet 1991).

6. Note that speech perception in this account already involves categorical features. Sub-featural properties in the speech signal are thus not considered relevant.

Fikkert and Levelt assume that all word forms in the lexicon compete for recognition. A representation is excluded as a possible candidate, if one of the perceived features mismatches the representation of that word. For instance, a lexical representation with the feature *labial* (e.g. from learning a word like /bɔn/) will be excluded, if the infant perceives a feature coronal (e.g. in the word /dɪn/). A representation is maintained as a candidate, if the perceived and stored forms *match* (i.e. are identical), or have *no mismatch*. Perceived and stored forms have no mismatch if there is no PoA feature specified in the lexicon (e.g. when learning /dɪn/), because no perceived feature can mismatch an underspecified (Ø) PoA. Table 2 shows how the perceived forms of four of the novel words used in the experiments reported in Fikkert (to appear) compare to their respective stored representations for Dutch.

Table 2. FUL proposal for the comparison between stored lexical representations and perceived forms of the same words, adapted from Fikkert (to appear)

Learned word	Stored representation	Perceived form	Comparison
/dɔn/	<i>labial</i>	coronal labial	Mismatch
/bɔn/	<i>labial</i>	labial labial	Match
/dɪn /	Ø	coronal coronal	No mismatch
/bɪn/	Ø	labial coronal	No mismatch

Fikkert (to appear) uses this application of the FUL model to predict which switches between which of the four words in Table 2 Dutch children will notice in the simple version of the switch paradigm, where only one word is presented during habituation. Her predictions are based on the comparison of the matching, as described above, between the same trials and the switch trials. When the matching differs between the switch trials and same trials, i.e. one is a match and the other is a mismatch, or one is a no mismatch and the other is a mismatch, it is predicted that children will recover from habituation and notice the switch.⁷ In this interpretation of the switch paradigm children thus habituate to the combination of an object, a phonological representation and a matching condition.

Table 3 summarizes Fikkert's predictions, based on the comparison between matching patterns, for infants' performance in the switch paradigm with the contrasts /bɪn/-/dɪn/

7. Note that this prediction applies irrespective of whether the same or the switch trial yields the mismatch. Therefore, it is not simply a question of whether the perceived form of the switch trial mismatches the stored representation because even on the same trials the child can experience a mismatch (e.g. /dɔn/, see Table 2).

and /bɔn/-/dɔn/. She predicts that 14-month-old infants will be able to detect a switch between /bɔn/ and /dɔn/, but not between /dɪn/ and /bɪn/.

Table 3. Predictions from Fikkert (to appear) for 14-month-olds’ performance in the switch procedure with the contrasts /bɪn/-/dɔn/ and /bɪn/-/dɪn/

Learned word	Stored representation	Trial	Perceived form	Comparison	Switch detection
/dɔn/	labial	same (dɔn)	coronal, labial	Mismatch	YES
		switch (bɔn)	labial, labial	Match	
/bɔn/	labial	same (bɔn)	labial, labial	Match	YES
		switch (dɔn)	coronal, labial	Mismatch	
/dɪn/	Ø	same (dɪn)	coronal, coronal	No mismatch	NO
		switch (bɪn)	labial, coronal	No mismatch	
/bɪn/	Ø	same (bɪn)	labial, coronal	No mismatch	NO
		switch (dɪn)	coronal, coronal	No mismatch	

With respect to vowel contrasts, the model predicts a successful detection of a switch only when replacing the learned label /bɔn/ of an object by the new label /bɪn/ but not vice versa. This is because the perceived form of /bɔn/, i.e. coronal coronal, and that of /bɪn/, i.e. labial coronal, match and mismatch, respectively, the stored representation of the learned word /bɔn/, i.e. *labial*. That is, in this case, same and switch trials differ in their matching, and therefore a switch can be detected. In contrast, both the perceived forms of /bɔn/ and /bɪn/ do not mismatch the stored representation of /bɪn/, i.e. Ø. That is, the switch and same trials yield the same matching pattern between perception and stored representation, and consequently no switch can be detected. Regarding the words /dɪn/ and /dɔn/, both the perceived form of /dɪn/ and that of /dɔn/ yield a mismatch with the representation of a learned word /dɔn/, i.e. *labial*, and they yield a no mismatch with the representation of a learned word /dɪn/, i.e. Ø. This means that when learning either /dɪn/ or /dɔn/, both the same and switch trials yield the same matching pattern, and therefore, it is predicted that no switch will be detected.

The results presented in Fikkert (to appear), which were reviewed in Section 2 above, confirm the FUL model’s predictions for words which minimally differ in both consonants and vowels. This success suggests that the model may be able to explain

other cases of early word recognition, which demonstrate a differential performance across minimally different word pairs. For instance, it could explain the noticing of the switch between /dit/ and /dɪt/ but not between /dit/ and /dut/ or /dit/ and /dʊt/ found by Curtin et al. (2009). We attempt here an application of the FUL model to the learning of these three vowel pairs. The English vowels /i/ and /ɪ/ have the same place of articulation, i.e. coronal, whereas /u/ is dorsal.⁸ Following Levelt (1994), Fikkert and Levelt (2008) and Fikkert (to appear), we assume that the feature coronal is underspecified. However, the three English vowels also differ in vowel height, i.e. /i/ and /u/ are high vowels, while /ɪ/ is mid-high. Given that neither study explicitly assumes the underspecification of any vowel height feature, we cannot *a priori* assume that high or mid-high are specified or not.⁹ Therefore, in our application of the FUL model, we consider the three possible scenarios, i.e. that both high and mid-high are specified, that they are underspecified, or that one of them is underspecified. Table 4 shows the perceived forms and *possible* stored representations of the three words used by Curtin et al.

Table 4. FUL model proposal for the comparison between stored representations and perceived forms for Curtin et al.'s words. First line: height featured is specified, second line: height feature is underspecified

Learned word	Stored representation	Perceived form	Comparison (PoA-height)
/dit/	Ø, <i>high</i> Ø, Ø	Coronal, coronal, high	No mismatch-Match No mismatch-No mismatch
/dut/	<i>dorsal, high</i> dorsal, Ø	Coronal, dorsal, high	Mismatch-Match Mismatch, No mismatch
/dɪt/	Ø, <i>mid-high</i> Ø, Ø	Coronal, coronal, mid-high	No mismatch-Match No mismatch- No mismatch

Although Fikkert tested her predictions using the simple version of the Switch paradigm, in which only one word is learned, we believe that the FUL model proposal can also be applied to the version of the switch paradigm where two words are learned, which was used in Curtin et al.'s study. In order to apply Fikkert's proposal to the learning of two words, a number of assumptions need to be made. For instance, Fikkert seems to propose that, during the habituation phase, infants learn (1) the association

8. Within the FUL model, the level of specification of a feature depends on its contrastive nature in the specific language (Lahiri & Reetz 2002). This means, for instance, that *rounding* will not be specified in English while vowel height and PoA, which are used to contrast many vowels in this language, must be specified.

9. See Levelt (1994) for an analysis of vowel height features in language production data of Dutch children.

between an object and a stored representation, and (2) the matching pattern between the object's perceived form and its stored representation. In what follows, we interpret Fikkert's proposal as stating that these two types of knowledge can be learned and stored for one, two, or any number of object-label combinations separately. Although Fikkert's proposal considers only one phonological feature per word, i.e. PoA, Curtin et al.'s data involve two features, i.e. PoA and vowel height. We assume that infants store both PoA and vowel height features (either specified or underspecified) and that each feature yields a separate matching pattern, which is then stored. Given the possible representations shown in Table 4, which follow this interpretation of Fikkert's proposal, one can make predictions for switch detections. As in Fikkert (to appear), these should be based on how the matching patterns of same trials, i.e. where an object-label association which occurred during habituation is presented, compare to those of switch trials, where the label of a learned association is switched with another label. Since we assume that matching patterns are stored per learned word, during the test phase, infants will expect a specific matching pattern for each object.

Following these assumptions, in the group of infants who were habituated to /dit/ and /dɪt/, same trials yield the same pattern for both words, i.e. no mismatch for PoA and match for height, if high and mid-high are specified in their lexicon, as can be read off Table 4. Both words yield no mismatch-no mismatch, if the vowel height features are underspecified. Recall that, according to Fikkert's application of the FUL model, a match or a no mismatch in either same or switch trials together with a mismatch in the other test trial leads to the noticing of the switch. Therefore, in order for babies to detect a switch between /dit/ and /dɪt/ in both directions, as was found by Curtin et al., switch trials should lead to a mismatch for vowel height.¹⁰ This is only the case if their stored representations are specified for vowel height. Thus, the model can only make a correct prediction for the recognition of /dit/ and /dɪt/, if the full specification of *high* and *mid-high* is assumed. Importantly, the same assumption yields to the correct prediction that infants should not be able to notice the switch between /dit/ and /dʊt/ in either direction, which is what was found in Curtin et al. The explanation underlying this prediction is that same and switch trials for the perceived forms of /dit/ and /dʊt/ yield the same matching patterns, i.e. no mismatch-match, when compared to the stored representation of /dit/, and mismatch-match, when compared to the stored representation of /dʊt/.

With the same vowel height specification assumption, the model predicts that infants should be able to notice the switch between /dɪt/ and /dʊt/ in both directions. This is because, in the case of /dɪt/, same trials yield a match for vowel height, i.e. the

10. Note that the vowels in /dit/ and /dɪt/ have the same PoA, and therefore there will not be a difference between their matching patterns for this feature in same and switch trials.

perceived form and the stored representation are both mid-high, while switch trials yield a mismatch for the same feature, i.e. the perceived form is high, while the stored representation is mid-high. Similarly, in the case of /dut/, the same trials yield a match for vowel height, while the switch trials yield a mismatch for the same feature. These match and mismatch patterns in the same and switch trials, respectively, predict the detection of a switch. However, Curtin et al.'s results show that infants could not detect the switch between /dit/ and /dut/ in either of the two directions. Therefore, with our interpretation of the FUL model, which states that per learned word the infant stores features and their matching patterns separately, the FUL model is not able to fully explain cases such as the ones presented in Curtin et al. However, it is necessary to explore all possible interpretations of this model before definitely rulling out its explanatory power. These explorations should focus on making the model more explicit in terms of dealing with multiple phonological features and their matching patterns simultaneously. In the next section, we describe two other proposals which also aim at explaining early word recognition performance.

2.2 Auditory-phonetic explanation

As mentioned in the previous section, although the application of the FUL model to early word learning can explain the difference in performance between /dit/-/dit/ and /dit/-/dut/, it is not able to explain why 15 month-old infants cannot notice the switch between /dit/ and /dut/. Curtin et al. explain the difference between the infants' performance on the three different contrasts by means of the phonetic cues involved. An acoustic analysis of the 10 tokens of each of the words /dit/ and /dut/, which were presented to the infants, showed that their vowels had a large difference on their first formant (F1) whereas the 10 tokens of the words involved in the contrasts /dit/-/dut/ and /dit/-/dut/ overlapped to some extent on this dimension. Apparently, non-overlapping F1 values seem to be a strong cue for infants to learn words which minimally differ in their vowels. In contrast, the second formant (F2) seems to be a weaker cue: Although the pairs /dit/-/dut/ and /dit/-/dut/ were completely separated on the F2-dimension, this did not help the infants to distinguish the words in the switch trial. In addition, the same separation was found for all three contrasts along the third format dimension, which is the acoustic correlate of the phonological feature rounding. This did not help infants to notice the switch between either /dit/-/dut/ or /dit/-/dut/, which suggests that this dimension is not exploited by infants of this age. This last finding is compatible with the underspecified view of the FUL model because the feature rounding is considered underspecified in English (Lahiri & Reetz 2002). In summary, non-overlapping F1 values, but not non-overlapping F2 or F3 values, seem to ensure infants' success at noticing the switch.

The difference in importance between F1 and F2 as cues to sound contrasts is also found in speech discrimination studies which do not involve word recognition.

That is, it has been shown that young infants better discriminate vowels along the F1 dimension than along the F2 dimension (Lacerda 1993; 1994). This initial bias in speech perception seems to shape infants' use of the phonetic cues in word learning. This is compatible with the PRIMIR view on language learning (Werker & Curtin 2005), in which it is proposed that several representational planes are involved in speech perception and word recognition. The raw acoustic cues are represented on a different level than the phonemes and the word representations. When a large cognitive load is devoted to one plane, such as the word plane, only the information of other planes which is easy to access can be used. Curtin et al. suggest several reasons why F1 would be more accessible to infants, including the higher acoustic energy in F1 and the greater stability of F1 across vowel productions.

The results from Curtin et al. also suggest that 15-month-old infants are able to exploit different sources of fine auditory/acoustic information separately when encoding word-object associations. This may mean that infants have only formed connections between their perception of F1 properties and their stored representations of words containing vowels, but not between their perception of F2 properties and their stored representations. Alternatively, one could think that both of these connections are formed, but that infants are only able to exploit one of them at a time in a word-recognition task. This is a developmental stage explicitly proposed within the Linguistic Perception (LP) model first put forward by Boersma, Escudero and Hayes (2003) and extended by Escudero (2005).

Both PRIMIR and LP propose that the inability to use F2 in word recognition is due to the fact that, in the absence of phoneme-like sound categories, infants will only use salient or most reliable auditory dimensions. In addition, PRIMIR suggests that a lowering of task demands may enable children to use less accessible phonetic information in word recognition, because children store every aspect of the word form as an exemplar. Contrary to PRIMIR, the LP model explicitly proposes that young infants can be in a developmental stage in which speech processing is uni-dimensional. That is, in early stages of language acquisition, infants process and store every auditory dimension, but separately. It is important to note that although the LP model implies that the word representations of infants are different from those of adults, it does not adhere to a discontinuity between speech perception and word representations. Rather, the model proposes that infants' use of auditory/acoustic information in speech perception is different from that of adults, which results in different word representations. In this respect, both the PRIMIR framework and the LP model are different from the FUL model, where discontinuity is assumed.

In the next section, we review studies which suggest that infants' failure to detect the switch may be due to the fact that the switch procedure is a very demanding task. If using methodologies which involve a lower cognitive load on the infants

resulted in a more successful performance than what was described until now, theoretical explanations of early word recognition should be modified to account for this.

3. Word recognition tasks with lower cognitive demands

Another explanation for the absence of a switch effect in most of the studies discussed in Section 2 is that the processing demands involved in associating an object with a label make it difficult for infants to pay attention to fine phonetic detail (e.g. Werker et al. 2002; Fennel & Werker 2003). In order to test this alternative hypothesis, Fennel has conducted a series of studies in which the demands of the switch procedure were lowered. One way of lowering the task demands is to present infants with less novel information during the learning phase. For example, if the object used for teaching the label-object association is already familiar to the infant, that aspect of the word learning task should require less processing demands. This was tested in a study reported in Werker and Fennel (2008). The parents of one participant group received a novel object 6 to 8 weeks prior to the test session. During these weeks, the infants got to know the object but, crucially, without a label. In the simple version of the switch task, the infants that were familiarized with the object prior to participating in the test did notice the switch from /dɪn/ to /ɡɪn/. They differed in this respect from the control group without prior familiarity with the object, as well as from all groups of infants in the studies discussed in the previous section.

A second potentially demanding aspect of the switch procedure is that the words are presented to the infants in isolation. The absence of any referential context might have made it more difficult for the children to understand that the sound string they were hearing was the label for the object with which they were presented. In order to test this hypothesis, Fennel and colleagues performed two manipulations. In the first, both the habituation and the test trials consisted of the target word embedded in the naming phrase “look at the...!”, thereby clearly establishing the referential link between the object and the label (Fennel 2006). In the simple version of the switch design, with only one word being learned, 14-month-olds did indeed detect the switch between /brɪn/ and /dɪn/. In a second manipulation (Fennel, Waxman & Weisleder, 2007), the referential context was established by teaching the children that they were involved in a naming game. To this end, children first saw three trials in which a familiar object (e.g. a shoe) was named with its (isolated) familiar label. Following these three trials, the habituation phase for the simple version of the switch procedure started. Again, a well-established referential context allowed the children to detect the switch. These studies have led Fennel, Werker and colleagues to believe

that infants can use their perceptual skills in word learning when they have processing resources available.

As can be expected, this is not an all-or-nothing matter because when referential context is established by a naming phrase and infants have to learn two minimally different words, as in the traditional switch procedure, they fail again (Fennell 2006). The conclusion thus must be that learning two minimally different words is in itself a cognitively demanding task as well (see also Swingley & Aslin 2007). Importantly, the findings that the reduction of task demands allows infants to notice the switch between /bɪn/ and /dɪn/ challenges Fikkert's application of the FUL model, because it predicts that infants should not be able to notice the switch between these two words, given that the feature coronal is not specified in the lexicon. It is not clear how exactly the LP model and the PRIMIR framework would explain the different results for the same contrast using different methodologies. What is important to consider is that in both proposals there are no underspecifications or global representations, and in principle, all features or auditory cues are represented and available for infants to use. In addition, within the PRIMIR framework, it is specifically proposed that task demands function as a filter for infants' attention to certain aspects of the speech signal.

Despite the influence of cognitive demands on early word recognition, it is important to remember that Curtin et al.'s findings show that infants can learn a minimal pair if the words differ in an important phonetic dimension. This was found even in the absence of a clearly established referential context. Thus, an explanation of early word recognition should also account for the variation in the type of phonetic detail that infants are able to utilize in word recognition, e.g. F1 vs. F2, voicing vs. place of articulation. This should be easily handled by the LP model and the PRIMIR framework. However, the models do not provide explicit accounts of cue-weighting for different sound contrasts.

Finally, when discussing infants' failure to notice the switch, we should consider the nature of the test phase. In that respect, Swingley and Aslin (2002) suggest that during the switch trials, the label that the infants hear partially activates the highly similar correct label associated with the displayed picture. This partial activation could inhibit the *novelty effect* (i.e. longer looking times) connected to the noticing of a switch, when the infant is presented with an object and an incorrect but closely related label. Yoshida, Fennell, Swingley and Werker (2009) tested this hypothesis. In their experiment, they used the training phase from the switch procedure, i.e. habituation, but a different test phase. At test, Yoshida et al. auditorily presented their participants with either /bɪn/ or /dɪn/, and visually presented them with the object "din" and the object "bin" side by side. During the first part of the test phase, the infants indeed looked reliably longer at the correct object upon hearing /bɪn/ and /dɪn/. Although this effect disappeared in the second part of the test phase, it shows that infants can learn that /bɪn/ and /dɪn/

are labels for two different objects. Thus, we must conclude that, despite the failure in the testing phase of the switch procedure, infants are able to access and encode fine phonetic detail in word learning.

Apart from constituting an important challenge to theories assuming under-specification in children's representations, Yoshida et al.'s findings also show that when investigating infants' representations of novel words, it is important to choose the method which best enables infants to demonstrate their actual knowledge. In the next section, we review research which has been conducted using a paradigm similar to that used by Yoshida et al. This paradigm seems to be more sensitive than the switch procedure when investigating the phonetic detail in infants' word representations.

4. Research on early word recognition using the mispronunciation paradigm

The second main paradigm used to investigate early word recognition is the *mispronunciation* procedure. In this procedure, children are presented with two pictures (e.g. a baby and a ball) and hear the label of one of the pictures. Infants tend to look longer at the picture of the object that matches the label. To assess infants' representations, their looking time at the target object upon hearing a correct (e.g. baby) versus an incorrect pronunciation (e.g. vaby) of its label are compared. If infants look shorter upon hearing the mispronunciation than upon hearing the correct pronunciation, this shows that the mispronounced form does not match their representation of the object's label as well as the correct form does. This procedure has mostly been used to assess infants' representations of familiar words, i.e. words that the infants know prior to the testing time, more specifically word onsets (Swingley & Aslin 2000; 2002; Bailey & Plunkett 2002; Van der Feest 2007), onsets of second syllables in bisyllabic words (Swingley 2003), and nucleus vowels (Mani & Plunkett 2007; Swingley & Aslin 2000; 2002). In all of these studies, it has been shown that mispronunciations distort word recognition of familiar words.

With respect to the specific phonetic detail that infants use in word recognition, Van der Feest (2007) conducted a mispronunciation study with 24- and 20-month-old infants. She presented the children with pictures of familiar words starting with the stop consonants /t/, /d/, /p/ and /b/. In the mispronunciations of these words, either place of articulation (PoA) or voicing was changed. The mispronunciations were therefore small and the nature of the mispronunciations was clearly defined. Overall, Van der Feest found that mispronunciations in PoA lead to a more severe distortion of word recognition than mispronunciations in voicing. Moreover, the direction of the mispronunciation matters: children detected a PoA mispronunciation if a labial

consonant was replaced by a coronal consonant (e.g. /p/ is mispronounced as /t/), but not vice versa. Regarding voicing, an asymmetry was also found: children detected the mispronunciation only if a voiceless consonant was mispronounced as a voiced consonant (e.g. /p/ is mispronounced as /b/), but not vice versa. In addition, there was a developmental difference between the two age groups, with only the older children being affected by the voicing mispronunciations. Van der Feest successfully explains her data by applying the FUL model proposal, in which the PoA coronal is underspecified in young children's representations.

Mani and Plunkett (2007) have conducted a series of studies to measure whether infants of 15, 18 and 24 months of age are able to detect large and small mispronunciations of familiar words, with the mispronunciations involving consonant and vowel contrasts. The authors found a reliable effect of mispronunciation of a single phonological dimension in the *vowels* of familiar words at 18 and 24 months but not at 15 months. However, in the same study, they found that infants of all three ages noticed a mispronunciation of a single phonological feature in the onset *consonants* of familiar words.

Ballem and Plunkett (2005) is one of the first studies to teach 14-month-old infants two novel words and test their recognition with the mispronunciation procedure. They found that the infants looked systematically at the picture associated to a newly learned word when hearing its correct pronunciation, but not when hearing a version with a mispronounced onset consonant. Unfortunately, the difference in looking times between correct pronunciations and mispronunciations was not significant, so that firm conclusions cannot be drawn from these results. However, Mani and Plunkett (2008) found that both 14- and 18-month-old infants looked longer at the picture of the newly learned words when they heard their correct pronunciations than when they heard a mispronounced version. Crucially, the correct pronunciations and the mispronunciations differed only in their vowel, e.g. /mɒt/ versus /mɪt/. These findings are in line with Curtin et al.'s findings which show that infants younger than 17 months can utilize fine phonetic detail of vowels when recognizing novel words. However, the mispronunciations of the vowels involved a change in three dimensions, namely vowel height, backness and rounding, which correspond to an auditory difference larger than in any of the contrasts used in Curtin et al.

The mispronunciation study which is most comparable to Curtin et al.'s switch study is Mani, Coleman and Plunkett (2008), where it was tested whether a mispronunciation of the vowel in either vowel height, vowel backness or vowel roundness, would be detected for novel words by 18-month-old babies. Their findings show that infants at this age can detect mispronunciations of vowel height (F1) and vowel backness (PoA, F2) but not of rounding (F3). This finding is partially compatible with Curtin et al.'s findings because they also found that infants did not exploit rounding or F3 differences when recognizing words. This is also compatible with Fikkert's phonological model

because this model assumes the underspecification of phonological features which are not contrastive in a language, in this case, the feature rounding in English.

Despite the similarities, Mani et al.'s (2008) study and findings differed from those of Curtin et al. in crucial ways. Mani et al. find that 18-month-old infants have the same success when detecting vowel height and vowel backness mispronunciations, whereas Curtin et al. find that 15-month-old infants can only detect vowel height or F1 differences but not vowel backness or F2 differences. These divergent findings could have a developmental explanation, i.e. that, at 15 months, infants can only exploit the F1 dimension and, at 18, they can exploit both F1 and F2. Recall that also in the switch paradigm, infants over 17 months are able to detect minimal differences between difficult contrasts. However, it is difficult to compare the findings of the two studies since they were conducted using two different procedures. As discussed before, children seem to be able to use some information when presented with a mispronunciation task than when presented with a switch task. This means that it is yet to be shown whether younger populations exhibit the same ability to detect differences in vowel height and vowel backness in a mispronunciation task. In addition, the infants in Mani et al. did not learn novel words that formed a minimal pair and the novel words were presented in a small phrase. These factors might have also contributed to their infants' relative success when compared to those tested by Curtin et al.

In the next section, we discuss specific directions that we believe are needed in order to resolve the methodological and theoretical problems outlined above in future research.

5. Summary and further research

This paper reviews the findings of studies investigating the use of phonological and phonetic detail in early word learning and recognition. These studies have been conducted using two different paradigms, namely the switch and the mispronunciation procedures. Studies conducted using the switch procedure have first shown that infants younger than 17 months of age fail to detect minimal differences between words differing in their onset consonant, e.g. /bɪn/-/dɪn/. In order to account for these results and to predict infants' performance in minimal word pairs differing in consonants and vowels, Fikkert (to appear) applies a phonological model, namely the FUL (Featural Underspecified Lexion) model developed by Lahiri and Reetz (2002). This model assumes that infants do not represent words in an adult-like fashion, but rather have underspecified representations for particular features. The predictions of the model are borne out in all of the studies reviewed in Fikkert (to appear), which were conducted with infants younger than 17 months of age. However, it is necessary to consider more thoroughly what happens in the switch task and what causes infants'

noticing the switch before we can conclude whether this account can also explain the findings reported in Curtin et al. (2009), who tested the recognition of novel words containing three different vowel contrasts by 15-month-old infants. These latter findings can be explained by models such as the LP model and the PRIMIR framework, which assume that infants assign differential weights to the auditory information involved in the sounds differentiating words. More recent studies have manipulated factors such as word familiarity and referential information, obtaining success with infants younger than 17 months of age in the differentiation of minimally different words. These same words could not be distinguished by the infants in previous switch studies. Different findings in different tasks are problematic for the FUL model, but are explicitly addressed in the PRIMIR framework, which incorporates an interaction between available information and processing demands.

In contrast, infants have been more successful at showing sensitivity to phonetic detail in word learning and recognition when tested in the mispronunciation procedure than when tested in the switch procedure. With regards to newly learned words, which were especially difficult for 14- or 15-month-olds in the switch procedure, Ballem and Plunkett (2005) found a difference between correct and mispronounced items, which did not reach significance. Studies conducted by Mani and colleagues showed that 14-month-olds' performance is affected by mispronunciations of the vowels in newly learned words, when the mispronunciations differ in three phonological features from the correct pronunciation. Other studies have also shown that mispronunciations in vowel height and vowel backness, but not rounding, affect infants' recognition. These results suggest that infants treat dimensions differently but do not replicate Curtin et al.'s results where infants succeeded at detecting a switch in vowel height, but not in vowel backness.

Further empirical research should enable us to fully understand the nature of children's early word representations as well as the processes involved in word learning and word retrieval or recognition. This type of research should be informed by theoretical proposals and thus test specific but contradicting hypotheses. Given the emphasis that theories such as the PRIMIR framework and LP model place on the role of phonetic cues, as opposed to phonological features, in word recognition, it is of great importance that the auditory properties of the stimuli are known. Furthermore, it should be made explicit that these properties give answer to a specific research question. Thus, for instance, if one aims at testing the comparative use of F1 and F2 in early word recognition, the vowels' F1 and F2 should be controlled in such a way that, for instance, /dit/ and /dit/ have almost equal F2 values but completely non-overlapping F1 values, while /dit/ and /dut/ have almost equal F1 values but completely non-overlapping F2 values. This methodological control of the stimuli has been adopted by Benders, Curtin, Escudero and Swingley (in progress), for auditory dimensions, and by Mani et al. (2008) and Van der Feest (2007), for phonological features.

According to the results of the studies reviewed in Sections 2, 3 and 4, the mispronunciation procedure represents the most sensitive method for examining the contents of early word representations. Thus, we suggest that this and similar methods should be preferred in future research to test theories about the contents of infants' representations. In order to examine the relative attention to specific phonetic detail, an adaptation to the original mispronunciation procedure can be considered. Within this adaptation, the minimal pair mispronunciation procedure, infants are taught two novel words which differ in two dimensions, e.g. /dit/ and /dut/, which differ in F1 and F2, or /bɪn/ and /tɪn/, which differ in place of articulation and voicing. Subsequently, two pictures of the newly learned words are presented side by side together with either a correct pronunciation or a mispronunciation. In this case, each mispronunciation is a mispronunciation of both words but in different dimensions. The examination of which object infants look at upon hearing the mispronunciation can reveal which dimension is most important in infants' lexical representations. For instance, if a child has learned the pair /dit/-/dut/, the new word /dɪt/ is a mispronunciation in F1 as well as a match in F2 for /dit/. On the other hand, /dɪt/ is a mispronunciation in F2 as well as a match in F1 for /dut/. If infants indeed exploit the F1 dimension to a larger extent than the F2 dimension, they will look at the object whose label matches the label they hear in F1. If they also exploit the F2 dimension, they will exhibit a shorter looking time to the F1 match upon hearing a mispronunciation.

The minimal pair mispronunciation methodology is similar to that used in Yoshida et al. and is comparable to the traditional switch task because infants have to learn a minimal pair. In addition, this methodology can provide a stringent test for models which aim at explaining the relative importance of auditory dimensions in early word learning, such as the Linguistic Perception model (LP, Escudero 2005; Boersma, Escudero & Hayes 2003), the PRIMIR framework (Werker & Curtin 2005), and Fikkert's (to appear) application of the Featural Underspecified (FUL) model. Recall that both the PRIMIR framework and the LP model explain Curtin et al.'s (2009) results by proposing that infants weigh F1 cues more heavily than F2 cues in vowel perception and early word recognition. If, for instance, it is found that infants detect only F1 mispronunciations and not F2 mispronunciations in the minimal pair mispronunciation procedure, both theoretical stands will be supported. Alternatively, if it is found that both F1 and F2 mispronunciations are detected, the PRIMIR framework, which predicts that less accessible cues may be used in less demanding tasks, will also be supported. Additionally, if a study shows that 15-month-olds use both F1 and F2 when tested using the suggested methodological procedure, it might be the case that these infants could potentially detect the switches between both /dit/ and /dut/ and /dɪt/ and /dʊt/, which were not detected in Curtin et al.'s study. This finding would again contradict the FUL model, which predicts success for /dɪt/ and /dʊt/ but failure for /dit/ and /dut/.

One important issue rarely addressed by previous studies is cross-linguistic variation and the acquisition of two languages at the same time, i.e. bilingualism. As mentioned in the introduction of this paper, it is well-known that speech perception skills are language-specific, i.e. they are optimized for the ambient language. Crucially, these language-specific abilities develop in the first year of life, before the age in which infants are commonly tested on their word recognition skills. If the development of speech perception indeed guides infants' word representations, cross-linguistic differences in the information used by infants are expected. Dietrich, Swingley and Werker (2007) compared Dutch and English 18-month-olds' interpretation of vowel duration as a cue to word learning. Their results show that only the Dutch infants, who are learning a language in which vowel duration is a signal for phonemic contrasts, and not the English children, learning a language where vowel duration does not signal phonemic contrasts, noticed the switch when only the duration of the vowel segment was changed. In addition, the cross-linguistic differences in the production of dimensions such as F1 and F2 for vowel categories may lead to differences in word recognition. For instance, Dutch exploits the F2 dimension to a larger extent than English does, and therefore it would be expected that Dutch infants could exploit this dimension earlier than English infants in a word recognition task. Regarding bilingual infants, Fennell and colleagues, in Ottawa, as well as Bosch and colleagues, in Barcelona, are currently conducting research on extending the monolingual studies to the bilinguals' first words.

Another issue which remains unsolved is the specific interrelation between the categorization of sounds and the use of such categories for purposes of word learning and word recognition. The categorization of speech sounds (ultimately) requires attention to and integration of several phonetic cues (Repp 1984). Thus, research should focus on understanding what types of information infants rely on and the relative attention they pay to these different dimensions. This means that more research similar to Van der Feest (2007, for stop consonants) and Mani, Coleman and Plunkett (2008, for vowels), who compared the effects of mispronunciations in different phonological dimensions, is very much needed. Although infants' speech discrimination skills are in place by the first year of life, the categorization of speech sounds is in development up to the primary school age (e.g. Gerrits 2001, for Dutch). This also means that we can expect their word representations to be different as well. This means that success in the mispronunciation procedure does not necessarily show that infants' word representations are completely adult-like. Thus, comparisons between different age groups are important, especially because a common assumption in all infant research on phonetic detail in word representations is that adults store all of this detail and have access to it in word recognition. For this reason, Van der Feest (2007) tested an adult control group on a subset of her stimuli and found that, for the adults, the effect of mispronunciation in place was stronger than the effect of mispronunciation in voicing. This mirrored her infant results and shows that if infants treat cues or features differently, this must

not necessarily be ascribed to a developmental leap. With regard to F1 and F2, to our knowledge, there are no studies which investigate adults' access to these cues in word recognition, which complicates the interpretation of Curtin et al.'s results with infants. If the computational demands affect infants' differential access to F1 and F2 in word recognition, we would expect not to find a difference between F1 and F2 in adults. However, psychoacoustic factors could apply to adults as well.

In sum, the study of age effects in word recognition and a thorough examination of the interaction between speech perception and word recognition could only be achieved if we can compare the behaviour of infants, young and older children, adolescents, adults, the elderly, and native and non-native listeners, using the same methodology. For instance, the minimal pair mispronunciation procedure introduced above compares well to that used in current psycholinguistic research on word recognition in native adult listeners (Weber & Cutler 2004; Shatzman & McQueen 2006), and L2 learners (Weber & Cutler 2004; Cutler, Weber & Otake 2006; Escudero, Hayes-Harb & Mitterer 2008; and Escudero, Broersma and Simon under review), as well as in research on speech sound categorization in infants (McMurray & Aslin 2004), and native as well as non-native adults (Flege 1995; Best 1995; Escudero & Boersma 2004). Further attempts following this line of research will be enlightening.

References

- Aslin, R.N., Jusczyk, P.W. & Pisoni, D.B. 1998. Speech and auditory processing during infancy: Constraints on and precursors to language. In *Handbook of Child Psychology: Cognition, Perception, and Language*, Vol. 2, D. Kuhn & R. Siegler (eds), 147–254. New York NY: Wiley.
- Bailey, T.M. & Plunkett, K. 2002. Phonological specificity in early words. *Cognitive Development* 17: 1265–1282.
- Ballem, K.D. & Plunkett, K. 2005. Phonological specificity in children at 1;2. *Journal of Child Language* 32: 159–173.
- Benders, T., Curtin, S., Escudero, P. & Swingle, D. In progress. Do infants use their phonetic skills in word learning? Infants' representation of F1 and F2 in novel words.
- Best, C.T. 1995. A direct realist view of cross-language speech perception. In *Speech Perception and Linguistic Experience: Theoretical and Methodological Issues*, W. Strange (ed.), 171–203. Baltimore MD: York Press.
- Boersma, P., Escudero, P. & Hayes, R. 2003. Learning abstract phonological from auditory phonetic categories: An integrated model for the acquisition of language-specific sound categories. In *Proceedings of the 15th International Congress of Phonetic Sciences, Barcelona*, M.J. Sole, D. Recansens & Romero, J. (eds), 1013–1016. Barcelona: Universitat Autònoma de Barcelona.
- Curtin, S., Fennell, C.T. & Escudero, P. 2009. Weighting of vowel cues explains patterns of word-object associative learning. *Developmental Science* 12(5): 725–731.
- Cutler, A., Weber, A. & Otake, T. 2006. Asymmetric mapping from phonetic to lexical representations in second-language listening. *Journal of Phonetics* 34: 269–284.

- Dietrich, C., Swingle, D. & Werker, J.F. 2007. Native language governs interpretation of salient speech sound differences at 18 months. *Proceedings of the National Academy of Sciences of the USA* 104(41): 16027–16031.
- Escudero, P. 2005. Linguistic Perception and Second Language Acquisition: Explaining the Attainment of Optimal Phonological Categorization [LOT Dissertation Series 113], Ph.D. dissertation, Utrecht University.
- Escudero, P. & Boersma, P. 2004. Bridging the gap between L2 speech perception research and phonological theory. *Studies in Second Language Acquisition* 26: 551–585.
- Escudero, P., Broersma, M. & Simon, E. under review. Recognition of auditorily confusable words in native listeners versus L2 and L3 learners.
- Escudero, P., Hayes-Harb, R. & Mitterer, H. 2008. Novel second-language words and asymmetric lexical access. *Journal of Phonetics* 36(2): 345–360.
- Feest, S. van der. 2007. Building a Phonological Lexicon: The Acquisition of the Dutch Voicing Contrast in Perception and Production. Ph.D. dissertation, Radboud University Nijmegen.
- Fennell, C.T. 2006. Infants of 14 months use phonetic detail in novel words embedded in naming phrases. *Proceedings of the 30th Annual Boston University Conference on Language Development, Boston, Massachusetts*, 178–189. Somerville MA: Cascadilla.
- Fennell, C.T. & Waxman, S.R. & Weisleder, A. 2007. With referential cues, infants successfully use phonetic detail in word learning. In *Proceedings of the 30th Annual Boston University Conference on Language Development*, Boston, Massachusetts, 206–217. Somerville MA: Cascadilla.
- Fennell, C.T. & Werker, J.F. 2003. Early word learners' ability to access phonetic detail in well-known words. *Language and Speech* 46(2): 245–264.
- Fikkert, P. To appear. Developing representations and the emergence of phonology: Evidence from perception and production. *Laboratory Phonology* 10.
- Fikkert, P. & Levelt, C. 2008. How does place fall into place? The lexicon and emergent constraints in the developing phonological grammar. In *Contrast in Phonology: Perception and Acquisition*, P. Avery, B.E. Dresher & K. Rice (eds), 231–268. Berlin: Mouton.
- Flege, J.E. 1995. Second-language speech learning: Findings, and problems. *Speech Perception and Linguistic Experience: Theoretical and Methodological Issues*, W. Strange (ed.), 233–273. Timonium MD: York Press.
- Gerrits, E. 2001. The Categorisation of Speech Sounds by Adults and Children. A Study of the Categorical Perception Hypothesis and the Developmental Weighting of Acoustic Speech Cues. Ph.D. dissertation, Utrecht University.
- Kuhl, P.K., Stevens, E., Hayashi, A., Deguchi, T., Kiritani, S. & Iverson, P. 2006. Infants show a facilitation effect for native language phonetic perception between 6 and 12 months. *Developmental Science* 9(2): 13–21.
- Kuhl, P.K., Williams, K.A., Lacerda, F., Stevens, K.N. & Lindblom, B. 1992. Linguistic experience alters phonetic perception in infants by 6 months of age. *Science* 255: 606–608.
- Lacerda, F. 1993. Sonority contrasts dominate young infants' vowel perception. *Journal of the Acoustical Society of America* 93: 2372.
- Lacerda, F. 1994. The asymmetric structure of the infant's perceptual vowel space. *Journal of the Acoustical Society of America* 95: 3016.
- Lahiri, A. & Reetz, H. 2002. Underspecified recognition. In *Laboratory Phonology* 7, C. Gussenhoven, N. Warner & T. Rietveld (eds), 637–676. Berlin: Mouton.
- Levelt, C.C. 1994. *On the Acquisition of a Place* [Doctoral Dissertation 8, Holland Institute of Generative Linguistics (HIL), Leiden University]. The Hague: HAG.

- Mani, N., Coleman, J. & Plunkett, K. 2008. Phonological specificity of vowel contrasts at 18-months. *Language and Speech* 51(1–2): 3–21.
- Mani, N. & Plunkett, K. 2007. Phonological specificity of vowels and consonants in early lexical representations. *Journal of Memory and Language* 57: 252–272.
- Mani, N. & Plunkett, K. 2008. 14-month-olds pay attention to vowels in novel words. *Developmental Science* 11(1): 53–59.
- McMurray, B. & Aslin, R.N. 2004. Anticipatory eye movements reveal infants' auditory and visual categories. *Infancy* 6(2): 203–229.
- Paradis, C. & Prunet, J.-F. (eds). 1991. *The Special Status of Coronals: Internal and External Evidence*. San Diego CA: Academic Press.
- Pater, J., Stager, C.L. & Werker, J.F. 2004. The lexical acquisition of phonological contrasts. *Language* 80(3): 361–379.
- Polka, L. & Werker, J.F. 1994. Developmental changes in perception of nonnative vowel contrasts. *Journal of Experimental Psychology: Human Perception and Performance* 20(2): 421–435.
- Repp, B.H. 1984. Categorical Perception: Issues, Methods, Findings. In *Speech and Language: Advances in Basic Research and Practice* 10, N. Lass, (ed.), 243–335.
- Shatzman, K.B. & McQueen, J.M. 2006. Prosodic knowledge affects the recognition of newly acquired words. *Psychological Science* 17(5): 372–377.
- Stager, C.L. & Werker, J.F. 1997. Infants listen for more phonetic detail in speech perception than in word learning tasks. *Nature* 388: 381–382.
- Swingle, D. 2003. Phonetic detail in the developing lexicon. *Language and Speech* 46(2–3): 265–294.
- Swingle, D. & Aslin, R.N. 2000. Spoken word recognition and lexical representation in very young children. *Cognition* 76: 147–166.
- Swingle, D. & Aslin, R.N. 2002. Lexical neighborhoods and the word-form representations of 14-month-olds. *Psychological Science* 13: 480–484.
- Swingle, D., & Aslin, R.N. 2007. Lexical competition in young children's word learning. *Cognitive Psychology* 54: 99–132.
- Weber, A. & Cutler, A. 2004. Lexical competition in non-native spoken-word recognition. *Journal of Memory and Language* 50: 1–25.
- Werker, J.F., Cohen, L.B., Lloyd, V., Casasola, M. & Stager, C.L. 1998. Acquisition of word-object associations by 14-month-old infants. *Developmental Psychology* 34(6): 1289–1309.
- Werker, J.F., Corcoran, K., Fennell, C.T. & Stager, C.L. 2002. Infants' ability to learn phonetically similar words: Effects of age and vocabulary size. *Infancy* 3: 1–30.
- Werker, J.F. & Curtin, S. 2005. PRIMIR: A developmental framework of infant speech processing. *Language Learning and Development* 1(2): 197–234.
- Werker, J.F., Gilbert, J.H.V., Humphrey, K. & Tees, R.C. 1981. Developmental aspects of cross-language speech perception. *Child Development* 52: 349–355.
- Werker, J.F. & Fennell, C.T. 2008. Infant speech perception and later language acquisition: Methodological underpinnings. In *Infant Pathways to Language: Methods, Models and Research Directions*, J. Colombo, P. McCardle & L. Freund (eds). Hove: Psychology press, 85–99.
- Werker, J.F. & Tees, R.C. 1984. Cross-language speech perception: Evidence for perceptual reorganization during the first year of life. *Infant Behavior and Development* 7: 49–63.
- Yoshida, K.A., Fennell, C.T., Swingle, D. & Werker, J.F. 2009. Fourteen-month-old infants learn similar-sounding words. *Developmental Science* 12(3): 412–418.

Restructuring head and argument in West-Germanic


Arnold E. Evers

The West-Germanic, c.q. Dutch, complex predicate appears as a series of verbs in predicate-final position. They were considered as a kind of separate unit (a field) in structuralist grammar (Bech 1955; Paardekooper 1955). A direct generative phrasing of that view follows from a head-to-head raising applied to a left-branching VP-stack consisting of the separate head-final predicates. This head-raising accounts easily for the strict adjacency of the verbs and for dialectal variations in their linear order. However, beyond its descriptive efficiency, the rule has little to recommend itself in present-day syntax, due to its properties of rightward movement, extracting heads and tugging heads in.

Yet, suppose there are SOV languages due to lexical frames with a predicate-final position of the V^0 . VPs stacked for a complex predicate in such environment would lack separately PF identifiable left edges to define their hierarchy and their subjects. Suppose this disqualifies the preservation of VP-stacks in SOV languages. The hierarchy of predicates is rather present in the selectional hierarchy of the (string adjacent) predicate final verbs. It is argued that V-to-V Raising is a variant of the V-to-I raising in SVO languages. It reconstructs the predicate hierarchy in a single complex head.

1. An intuitive view on the problem

By restructuring a head, I mean the licensing of a predicative head by adjoining it to the sister (governing) head outside of its projection. The French V^0 -to- I^0 movement is an example. By restructuring an argument, I mean the licensing of an argument outside the projection it has been selected by. Subject formation following a V^0 -to- I^0 movement is the best example. The two licensing procedures are closely related. The V^0 -to- I^0 movement creates an extended projection IP. That prepares the position Spec, I for the subject in (1).




The head and argument restructuring offered by the West-Germanic V(erbal)-cluster is a marvel, since the head and argument restructuring get into a recursion. See example (2a) for German and (2b) for Dutch.

- (2) a. als [der König]^{1,2} [seine Tochter]^{2,3} [ihren Frosch]^{3,4} [den Ball]^{4,5}
 [aus₅ dem Teich hochbringen₄ lassen₃ sehen₂ konnte₁]
 when [the king] [his daughter] [her frog] [the ball]
 [out of the pond dive-up let see could]
 when the king could see his daughter letting her frog dive up the ball out of
 the pond
- b. toen [de koning]^{1,2} [zijn dochter]^{2,3} [haar kikker]^{3,4} [de bal]^{4,5}
 [uit₅ de vijver kon₁ zien₂ laten₃ opduiken₄]
 when [the king] [his daughter] [her frog] [the ball]
 [out of the pond could see let dive-up]

The examples have all been taken from the same three-person fairy tale to facilitate their contextualization. There is at the right-hand side in (2) a row of subscripted verbs that select each other. Each of the verbs also selects a DP argument and these are lined up at the left-hand side of the construction in a structural case (one nominative, three accusatives). Superscripts on the DP arguments indicate an understood subject-predicate relation due to an extended V-V projection. The main direction of verb-verb selection in German is to the left {4-3-2-1} and the main direction in Dutch is to the right {1-2-3-4}, which yields the mirror effect between Dutch and German verb clusters. Within the standard languages themselves, and even more so in the dialect area, there are minor differences in linearization options. See for surveys Schmid and Vogel (2003), Wurmbrand (2001), Barbiers (2005). The Dutch and German structures are nevertheless close parallels. Fairly simple arguments suggest that the verbal rows in both languages form a single tight constituent, the V-cluster. The V-cluster constructions make the impression to be shape-conserving. They seem to return to the form of the simplex V⁰ predication. See the V⁰ structures in (3) and (4).

- (3) weil/omdat (because)
-
- DP DP DP PP V⁰+I⁰
- [die Prinzessin] [den Frosch] [den Ball] [aus dem Teich] [hochbringen₃ lassen₂ hat₁]
 [de prinses] [de kikker] [de bal] [uit de vijver] [had₁ laten₂ opduiken₃]
 the princess the frog the ball out the pond had let up-dive
 because the princess had let the frog dive up the ball out of the pond



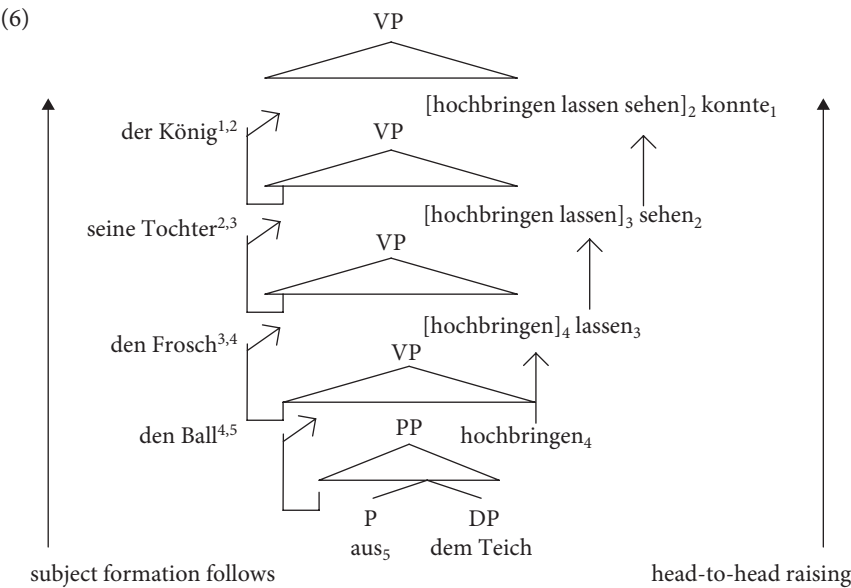
A parallel structure between simplex and complex predicates as in (3) and (4) will correctly predict that rules for argument distribution need not be complicated for the simplex and complex status of the predicate. Rules that get the new generalizing freedom over simplex and complex predicates are for example argument preposing, scope of negation, clitic movement, extraposition, reflexivization indefiniteness effects, and the assignment of neutral sentence stress. It is attractive to fuse lexical forms in such a way that their predicative heads V^0 are united into a complex head V^0 , while the arguments are attached to this head as in (3) parallel to (4). The deeper problem is how the theory of grammar can be invoked to turn a multiplicity of predications into a single predication. The transition from a hierarchical multi-clausal structure into a mono-clausal predication has two faces. It builds structure, the V-cluster, and it destroys structure, the separate predications.

It may be useful to add that the V-cluster cannot be seen as a morphological unit. Reasons, mainly taken from Haider (1993), are (i) V-cluster members are infinitives, *zu*-marked infinitives, past participles or finite verbs, whereas verbs in West-Germanic morphological compounds appear as verb stems only; (ii) there is a limited amount of order variation possible between the members of the V-cluster, but without semantic effects; (iii) except for a few idioms, real compounds cannot be a verb or have a verb as head (see Lieber 1983 for an explanation), nor (iv) can a compound member V^0 have a theta role-assigning effect outside of the compound, whereas this is the standard case in V-clusters; (v) the V-cluster allows sub-extraction of the finite verb in root clauses for the V-second effect. Like the internal order variation in (ii), this is incompatible with a morphological construct.

Let now head-restructuring and argument-restructuring be defined as in (1), then the generalization in (5) can be argued. The one requires the other.

- (5) West-Germanic
- | | | |
|--------------------|-------------------|------------------------|
| Head-restructuring | \Leftrightarrow | Argument-restructuring |
| V-cluster | | (subject formation) |

In West-Germanic all V-cluster constructions involve argument-restructuring and the other way around. All cases of argument-restructuring are based on a V-cluster that builds up an extended projection. It will be proposed that subject formation is the driving force behind this state of affairs. Successive V-to-V Raising creates each time an extended projection and thereby an external argument position. The external argument position is filled in by a structural case argument. See the diagram in (6).



The multiple subject formations in (6) follow the successive head restructurings. The head restructuring creates the extended projection needed for the subject as external argument. It remains still unclear how the multi-clausal structure in (6) turns at the same time into the mono-clausal argument space in (3).

1.1 Summary

Section 2 ('Subject Formation') deals with the structure-building aspect of V-to-V Raising (Section 2.1) and with the structure-destroying aspect (Section 2.2). The main point of Section 2.1 ('V-to-V Raising as projection extension') will be that the successive introduction of external arguments follows from a successive cyclic V-to-V Raising. See diagram (6) above. The transparency of the resulting structure can be achieved by claiming that the VP-hierarchy of empty places in a structure like (6) enjoys Government

Transparency as proposed in Baker (1988: 64). An alternative would be to drop the use of traces altogether. It will be argued in Section 2.2 ('Traces') that a grammar without traces is more strictly derivational and more straightforward about the relation between syntax and semantics.

Section 3 ('Syntactic Arrangement') will discuss a few descriptive problems. These are: (Section 3.1) the function of the infinitival prefix *zu/te*, (Section 3.2) the possibility of head movement, and (Section 3.3) the problem with long passives.

Finally, Section 4 ('Psychological Reality') will consider why all this distributional puzzling is supposed to be psychologically relevant.

2. Subject formation

2.1 V-to-V Raising as projection extension

Dutch/German infinitival projections sometimes remind of gerunds in English. See (7).

- (7) ... *den Ball aus dem Teich hochbringen ist gefährlich.*
 the ball out the pond up-dive is dangerous
 'Diving up the ball out of the pond is dangerous.'

The sentence in (7) does not specify who is going to make the dive. Yet, when the infinitive *lassen* ('let') is added as in (8), the subject of *hochbringen* ('dive up') may be specified.

- (8) *ihren Frosch^{3,4} den Ball^{4,5} [aus₅ dem Teich hochbringen₄ lassen₃], ist*
 her frog the ball out the pond up-dive let, is
 (letting her frog dive up the ball out of the pond, is)

The causative *lassen₃* does not further qualify the external argument (*ihren Frosch^{3,4}*) of *hochbringen₄*. It rather licenses its syntactic position. This type of grammatical licensing repeats itself in the example set (9) below. Each new structural argument is related to a combination of two predicative heads. One of the heads, like *hochbringen₄* in (8)/(9d), delivers the external theta role, while the other predicative head *lassen₃* delivers the external syntactic position. The same double-head analysis for the external argument applies as well to the small clause subject *den Ball^{4,5}* ('the ball') in (9a,b). The structural argument *den Ball^{4,5}* is the external argument of the result clause *aus dem Teich* ('out of the pond') and syntactically licensed by *hochbringen₄* ('dive up') in (9a,b).

- (9) a. *den Ball [aus₅ dem Teich] ... ist gefährlich
the ball out the pond ... is dangerous
- b. den Ball^{4,5} [aus₅ dem Teich hochbringen₄]...
the ball [out the pond up-dive] ...
- c. *ihren Frosch den Ball [aus₅ dem Teich hochbringen₄] ...
her frog the ball [out the pond up-dive] ...
- d. ihren Frosch^{3,4} den Ball^{4,5} [aus₅ dem Teich hochbringen₄ lassen₃] ...
her frog the ball [out the pond up-dive let] ...
- e. *seine Tochter ihren Frosch^{3,4} den Ball^{4,5} [aus₅ dem Teich hochbringen₄ lassen₃] ...
his daughter her frog the ball [out the pond up-dive let] ...
- f. seine Tochter^{2,3} ihren Frosch^{3,4} den Ball^{4,5} [aus₅ dem Teich hochbringen₄ lassen₃ sehen₂] ...
his daughter her frog the ball [out the pond up-dive let see] ...
- g. *der König seine Tochter^{2,3} ihren Frosch^{3,4} den Ball^{4,5} [aus₅ dem Teich hochbringen₄ lassen₃ sehen₂] ...
the king his daughter her frog the ball [out the pond up-dive let see] ...
- h. (weil) der König^{1,2} seine Tochter^{2,3} ihren Frosch^{3,4} den Ball^{4,5} [aus₅ dem Teich hochbringen₄ lassen₃ sehen₂ konnte₁] ...
(because) the king his daughter her frog the ball [out the pond up-dive let see could] ...

The new structural case that legitimately appears in (9b), (9d), (9f), (9h) gets a subject function. This is seen if one changes the last infinitive into a finite verb, compare (9d) with (10d), (9f) with (10f), (9h) with (10h).

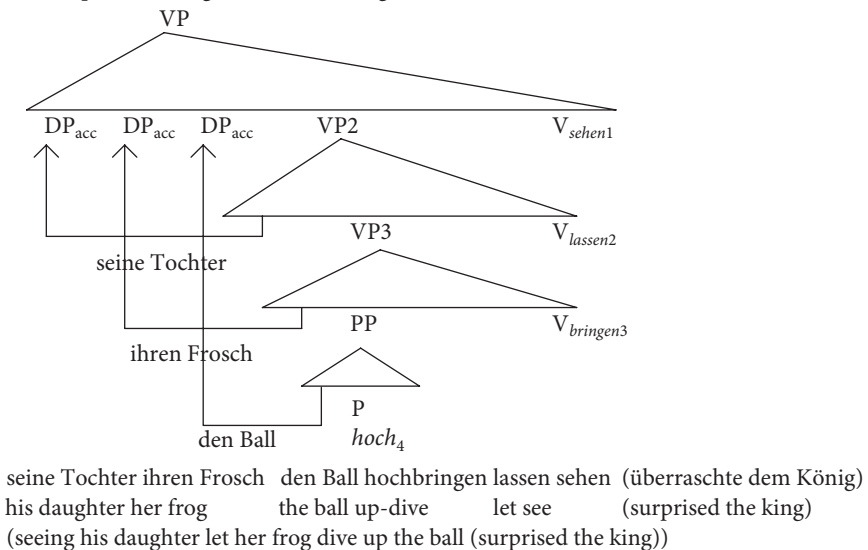
- (10) dass (that)
- d. ihr Frosch^{4/fin} den Ball^{4,5} [aus₅ dem Teich hochbringt_{4/fin}]
her frog the ball [out the pond up-dives]
- f. seine Tochter^{3/fin} ihren Frosch^{3,4} den Ball^{4,5} [aus₅ dem Teich hochbringen₄ lässt_{3/fin}]
his daughter her frog the ball out the pond up-dive lets
- h. der König^{2/fin} seine Tochter^{2,3} ihren Frosch^{3,4} den Ball^{4,5} [aus₅ dem Teich hochbringen₄ lassen₃ sieht_{2/fin}]
the king his daughter her frog the ball out the pond up-dive let sees
(the king sees his daughter letting her frog dive up the ball out of the pond)

The finite verbs in (10) stop a recursive stacking of subjects. It is not remarkable that adding a new predicative head adds the possibility of an additional structural argument. What merits the attention in (9) is the licensing of the new arguments. Each new predicative head assigns no more than a single structural case and thereby it licenses

the external argument of its sister. The structural case *den Ball*^{4,5} (9b), *ihren Frosch*^{3,4} (9d), *seine Tochter*^{2,3} (9f), *der König*^{1,2} (9h) are respectively due to the projection extensions by *hochbringen*₄ in (9b), *lassen*₃ in (9d), *sehen*₂ in (9f), and *könnte*₁ in (9h). The procedure fits Burzio's (1986) requirement that an accusative will not be assigned unless there is another external argument in the offing.¹ Alternatively, each structural case now indicates that an external argument has been licensed by the fusion of two predicative head factors. West-Germanic V-to-V Raising is thereby unmasked as the recursive switch for structural case configurations.

Two points of the present analysis were already made by Zwart (1996) and by Koopman and Szabolcsi (2000). First, they had their structural case arguments systematically represented as external arguments and second, they separated their structural arguments from a complex predicate constituent containing the predicative heads. For theoretical reasons, Zwart (1996) as well as Koopman and Szabolcsi (2000) avoid head movements. Their analyses start with a VP-stack. Their VP constituents select but do not license the structural case arguments. Structural specifier positions are postulated in a left periphery outside of the VP-stack. These specifier positions serve to license the structural arguments. Hence, their structural case arguments move out of the VP-stack for case-licensing. The structural arguments leap-frog over each other into these specifier positions, see (11).

(11) DP path-crossing for structural targets



1. In the present case, the thematic accusative *den Ball* derives from the (result) small clause *aus dem Teich*. That has a satisfactory consequence. Small clause predicates are ⟨-V⟩ and will not restructure a complement. By consequence, they always are the most embedded part of a

Their West-Germanic matrix structure must now provide as many accusative configurations as there are structural arguments selected in the VP-stack. Their structural arguments move up into the matrix structure (directly or by intermediate steps). Remarkably, the eventual linear order of the structural arguments will again reflect the selectional order of their VP-origins in a rigid way. The order of the predicative heads by contrast appears to be less rigid. Within each West-Germanic dialect, and even more so in cross-dialect comparison, all kind of order variation is possible (see Schmid & Vogel 2003; Wurmbrand 2001; Barbiers 2005). Hence, one arrives in a somewhat paradoxical situation for the linear order of elements. Their analysis in (11) by leftward movements is basically driven by case-licensing configurations. Yet, after a certain amount of rearrangements, the structural arguments end up in a rigid left/right copy of the original order. The predicative heads, by contrast, are not primary targets of movements. Yet, they show in cross-dialect comparison much variation in their linear line-up.

The approach proposed here in (6) succeeds to escape the paradox between the rigid argument-order versus less rigid head-order. Its driving force is subject formation by projection extension. Projection extension is brought about by V-to-V Raising. Different linear orders of predicative heads may now follow directly when V-to-V Raising varies between head-adjunction to the left (German) and head-adjunction to the right (Dutch). The strict linear order of the structural arguments, by contrast, follows directly from the successive cyclic subject formations. The unpleasant crossing of paths towards structural case targets, originating from Chomsky (1995), is avoided. No stipulation of multiple Object-Agr positions is needed, nor is any movement of structural case arguments called for. An appropriate delay of external argument merge suffices. The structural argument should not come in (by internal merge) until its structural position has been shaped by V-to-V Raising. The different semantic functions of the three accusatives in (11) follow in (6) from the sub-predicates they are attached to as external arguments. This is in (6) for *den Ball*^{3,4} ('the ball') the predicate *hoch*₄ extended by *bringen*₃, for *den Frosch*^{2,3} ('the frog') the predicate *bringen*₃ extended by *lassen*₂, and for *seine Tochter*^{1,2} ('his daughter') the predicate *lassen*₂ extended by *sehen*₁. The thematic theta role *den Ball*^{4,5} ('the ball') must appear as the most embedded structural argument, because it is theta-qualified by a non-verbal small clause particle *hoch*₄, inevitably the most embedded predicative head.

2.2 Traces

It used to be common for constituent movement to leave a trace, a syntactically labeled but phonologically empty position. Traces guarantee that movements may add structure

complex predicate. This in turn guarantees that the thematic argument that they bring in is bound to get the most embedded structural case. It seems very well possible to arrange a small clause analysis for all thematic arguments (Mulder 1992).

without destroying their point of origin. Traces are a representational device. Due to traces the last derived representation includes all the previous syntactic configurations. Baker (1988) uses them to show how polysynthetic verbs can be derived by assuming an underlying syntactic head-complement structure in which the complement is theta-related to the matrix head. The incorporation of the complement head into the matrix head successfully predicts a polysynthetic verb and its properties. In the first place, (due to restrictions on syntactic movement) head incorporation is possible from complement phrases, but not from subject or adjunct phrases. In the second place, the new complex head governs into its complements. Grammatical head-argument functions that define argument preposing, scope of negation, reflexivization, case assignment, or cliticization hold for the new complex head and arguments in the complement.

Baker (1988: 64) derives the mono-clausal effect of long passive and long nominalization or long clitic movement by stating a government transparency for the trace-headed complement.

- (12) [object (den Ball) t_V] $_{VP}$ [V V] $_V$ $_{VP}$
- \uparrow \uparrow
 government transparency

There is another way to reach the same mono-clausal outcome. Baker's results are less dependent on traces than it seems. The important point is the locality of movements. Suppose that a head movement must target the nearest head there is, without crossing a subject or adjunct clause boundary (since these are not theta-governed). This direct restriction on syntactic movement derives the same head clusters, but without reference to traces. Suppose now that syntactic movements after all leave no traces. Then the raising of the complement head in a structure like (13) will remove at the same time the projection VP. This is not a new stipulation, but rather an automatic effect of phrase structure. No head V, no projection VP.

- (13)
- \rightarrow
 Pruning
 (VP) and (t_i)

Derived mono-clausality, arguments in a single V-projection, now follows without the use of traces and government transparency. The arguments that were licensed by the projection-line that has now disappeared (been pruned) must be reconnected to the structure. I follow here the proposal made by Seuren (1996: 51) that the stranded constituents reconnect to the next VP-projection that contains them.

In a sense, the stranded arguments were successfully attached to the VP-projection they were in and that does not change. See the DP argument in (13). It continues to be connected to the first VP it is in. It is the V-VP projection line that changes into an extended projection.

One may feel that lack of traces (or structure copies) will block a semantic interpretation of the syntactic structure, but that is not true. Lack of traces rather forces one to make the semantic rules as derivational as the syntactic rules. A semantic representation, in which all relevant semantic structures are simultaneously present, is a cognitive result, a final achievement. Alternatively, semantic distinctions may be seen as a series of fleeting moments that light up in succession due to the ongoing syntactic procedures. Semantics surfs on the wave of syntactic restructuring. Syntactic argument merging and semantic theta assignment, syntactic subject formation and semantic predicate identification, like syntactic *wh*-movement and semantic scope-assignment, or like syntactic V-second and semantic illocution marking must now operate simultaneously and pair-wise. They become immediate examples of semantic syntax as in, for example, Categorical Grammar.² There remains of course a relation, say for *wh*-phrases, between a scope-identifying and an argument-identifying position. That relation is now expressed in the derivational mode rather than in a representational mode. Things like, say, a *that*-trace filter will simply return as (in)ability to recover an argument position.

3. Syntactic Arrangement

3.1 Control

German attributive constructions have a fairly elaborated agreement relation with the DP they qualify, see (14).

- (14) ein_i von dem Frosch aus dem Teich hochgebrachter _{i} Ball
 a by the frog out the pond up-dived ball
 (a ball dived up out of the pond by the frog)

2. The syntax/semantics bifurcation reminds of the horsepower in a machine. There is horsepower in a properly fuelled machine. Yet, through the ages metaphysicians ended up feeling that this will not entitle us to postulate that there are horses in the machine (Ryle 1947). Following that line here, we conclude to one grammatical entity and two ways to describe it, syntactic distributions versus semantic functions. Arranging for two cyclic rules, one in syntax ('move') and one in semantics ('co-index'), stops if the movement traces themselves disappear.

The attributive endings reflect more or less explicitly case, gender, number and indefiniteness. The resulting set of ϕ -values appears in the attributive ending that is identified later on with the ϕ -set values of the DP that is being qualified. The attributive ϕ -set is in a sense controlled by the DP that it qualifies. The extensive attributive agreement (case, gender, number, definiteness) makes it somewhat improbable to assume that the attributive relation will in addition require an empty non-specified PRO element to be related to the D that it qualifies. When Control is described as a direct relation between the ϕ -sets of the attributive A-projection and the element D^0 of the qualified DP, an intermediate element PRO is avoided. As long as the grammatical marking on the predicate associates with a strictly local but predicate external controller, PRO can be left out and grammatical wellformedness is reached.

This detour over attributive constructions has been made because there is also Control of infinitival complements and the standard description for that type of Control is in trouble. Controlled infinitival complements are traditionally assumed to contain an internal subject position. Some more construction-specific principle subsequently guarantees that the subject position will remain empty (PRO). Further, the PRO element will be obligatory bound when there is a theta-marked antecedent in the matrix construction. The relevant point of PRO in the present context is that the controlled infinitival complements display an internal satisfaction of the subject requirement for predicates (due to the element PRO). What is the problem now? As a matter of fact, some controlled infinitives in West-Germanic allow V-to-V Raising. If we continue to assume that V-to-V Raising is triggered for complements that lack an internal subject configuration, an automatic absence of PRO in controlled infinitival complements is more than welcome. The good news is now that we get rid of an empty place PRO and that V-to-V Raising for controlled infinitival complements remains an option. The other side of the story is that we are committed to make a different proposal for Control if no V-to-V Raising follows.

Two properties of controlled infinitival complements in West-Germanic stand out, see (15) and (16).

- (15) they are all without exception *zu*-marked.
- (16) they are all either extraposed (to the right of the matrix verb).
or they are not extraposed, but reconstructed by V-to-V Raising.

The point in (16) may also be formulated in the following way. There is for controlled infinitival complements (in West-Germanic) without exception an alternative between either extraposition or reconstruction by V-to-V Raising. The V-to-V Raising seems optional for controlled infinitival complements, because the extraposition is. See the exemplification in (17).

- (17) a. (underlying structure)
 omdat de kikker [de bal op te duiken]_{VP} probeerde
 because the frog [the ball up to dive] tried
- b. (extraposition)
 omdat de kikkerⁱ probeerde [de bal op teⁱ duiken]_{VPi}
 because the frog tried [the ball up tot dive]
- c. (restructuring)
 omdat de kikkerⁱ de bal op [probeerde teⁱ duiken]_V
 because the frog the ball up [tried to dive]
 (because the frog tried to dive up the ball)

The element *zu/te* can be seen as an operator that explicitly marks infinitival projections as lacking a subject configuration. I take this position over from Haider (1984) who explains thereby how a *zu/te* has an intransivizing effect as in (23) below, when *zu/te* is not properly governed. The *zu/te* operator is not properly governed in the underlying structure (17a) either. There are two ways to get *zu/te* governed. Both result in a configuration in which the *zu*-marked predicate is locally connected with an external reference. It is either an extraposition of the full clause as in (17b), or a V-to-V Raising as in (17c). In both cases the element *zu/te* becomes a factor in the matrix clause. I now assume (as in Evers 2003) that the *zu/te* element must be co-indexed with a theta-marked DP antecedent within the local domain (the matrix clause). This directly involves the infinitival prefix *zu/te* in getting the predicate head *duiken* ('dive') subject-identified, see the indices in (17b) and (17c). The direct involvement of *zu/te* with the Control relation avoids the element PRO, but all major properties of Control follow. (i) Due to the definition of *zu/te* as a subject operator, Control can only concern the subject function; (ii) Control is excluded in constructions without *zu/te*, i.e. finite clauses, small clauses and infinitival clauses without *zu/te*; (iii) Binding of controlled *zu/te* takes place in exactly the same locality as binding of the bound anaphor *zichzelf* ('himself'). The co-indexing of *de kikker* ('the frog') with *zu/te* in the extraposed infinitival complement (17b) takes place between major constituents in the cycle defined by *probeerde* ('tried').

If the *zu/te* complement does not extrapose, *zu/te* is not activated. The verbal complement lacks a subject and must restructure, as if no *zu/te* had been present. After the V-to-V Raising *zu* is V-governed, activated as in (17c) and ready to carry the external theta role of the *zu/te*-marked infinitive. The co-indexing of *zu/te* in the restructured (17c) takes again place within the projection of the matrix verb *probeerde* ('tried'). The fact that *zu/te* is impervious for agreement morphology tallies well with its reflexive function.

The relation between V-to-V Raising and *zu/te* marking may be characterized as follows. Infinitival complements without *zu/te* marking {causatives, perceptuals, modals and certain aspectuals} are without exception restructured by V-to-V Raising.

Infinitival complements that are selected with a *zu/te* marking are not restructured as far as they are extraposed. A successful extraposition of *zu/te* marked infinitival complements is (without exception) possible for all verbs of Control. If extraposition is not applied to controlled complements, they will restructure by V-to-V Raising as in (17c). The V-to-V Raising for controlled complements is optional only because extraposition can be applied optionally.³ There is a group of verbs that is sentence-qualifying in some modal or aspectual sense {*schijnen te/scheinen zu* ‘seem to’; *hoeven te/brauchen zu* ‘need to’; *beginnen te/beginnen zu* ‘begin to’; *plegen te/pflegen zu* ‘be in the habit of, usually happen to’}. If the *zu/te*-marked complements of these verbs are extraposed, the sentence-qualifying matrix verb will fail to offer a theta-marked controller and the derivation crashes. See (18a).⁴ The alternative, V-to-V Raising, yields the extended projection for the external theta-role in the usual way. See (18b).

- (18) a. *weil es/König Siegmund scheint [den Frosch zu verstehen]_{VP}
 because it/king Siegmund seems [the frog to understand]
 (because it/the king seems to recognize the frog)
- b. weil König Siegmund den Frosch [zu verstehen scheint]_V
 because king Siegmund the frog [to understand seems]
 (because king Siegmund seems to understand the frog)

The analysis in this section aims to explain why V-to-V Raising for controlled complements is possible, but optional, whereas V-to-V Raising is obligatory for *zu*-marked complements of sentence-qualifying verbs, as in (18). The element PRO is avoided without an obvious loss of significant generalizations.

3.2 Head movement and pied-piping


There have been several attempts to portray V-cluster constructions as phrasal rearrangements that make no reference to heads X^0 (Zwart 1996; Koopman & Szabolcsi 2000; among others). Koopman and Szabolcsi (2000) show how that might be done for Hungarian and Dutch V-clusters in a parallel fashion. Before they begin to spell out that approach, they formulate an interesting objection against a derivation of V-clusters by means of head movement Move V^0 (Koopman & Szabolcsi 2000: 25; Section 3.7). They state their objection for Hungarian constructions, but the objection has a general quality which allows me to restate it for Dutch.

3. The present approach also implies (incorrectly) that V-to-V Raising is an option for all verbs of Control. This is certainly an overgeneralization, as pointed out in Rutten (1992: 78). Some additional factor of *zu/te* is still to be figured out.

4. (18a) is grammatical and even preferred in certain dialects. Presumably these allow VP-Raising no matter whether the complement is or is not *zu*-marked, see below under (20).

Consider the Dutch construction (19).

(19) toen



C⁰ she the frog the ball out the pond had let up-dive
(then she had made the frog dive up the ball out of the pond)

- (i) Move V⁰ as V-tot-V raising; no pied-piping
- (ii) Move V⁰ as V-to-V Raising with partial pied-piping
- (iii) Move V⁰ as V-second excorporation; no pied-piping

(i) First, the most embedded verb *opduiken*₃ ('dive up') moves to the right of the causative *laten*₂ ('let'). This yields the extended projection [*laten*₂ *opduiken*₃] ('let dive up'), which allows the licensing of the external argument for *opduiken*₃ ('dive up'), the accusative *de kikker*^{2,3} ('the frog'), within VP₂. The movement of *opduiken*₃ ('dive up') strands its original context in VP₃ [*de bal uit de vijver*] ('the ball out of the pond'). (ii) The next Move V₀ raises the causative *laten*₂ ('let') to the auxiliary *had*₁ ('had'). It yields the extended projection [*had*₁ *laten*₂] ('had let') in VP₁ and allows the licensing of the external argument for the causative, the nominative *zij*^{1,2} ('she'). The second movement Move V⁰ strands its external argument *de kikker*^{2,3} ('the frog'), but it pied-pipes obligatorily the dependent verb *opduiken*₃ ('dive up'). (iii) The third movement Move V⁰ brings the finite verb *had* ('had') into the position of the illocutionary operator for root clauses. This movement strands all arguments and even strands all former cluster members. It excorporates out of the cluster.

It is clear that Move V⁰ cannot operate without a few plausible principles that switch off pied-piping (first Move V⁰), switch it on (second Move V⁰), and switch it off again (third Move V⁰, the V2nd rule). There often are good reasons for stranding licensed context or for pied-piping it. For example, stranding the former context of *opduiken*₃ ('dive up') in VP₃ [*de bal uit de vijver* t₃] ('the ball out of the pond t₃') is justified if one accepts that the stranded material is reattached or re-licensed in VP₂, as argued in Section 2.2. The non-stranding but pied-piping of V₃ when its companion V₂ was raised to V₁, seems rational. The V₃ cannot be re-licensed as a subordinate predicate head. I have no proposal for the non-Pied-Piping behavior of the V2nd rule, the third example of Move V⁰. In general, it may be said that all rules may have the format Move X⁰ and that phrasal rules Move XP appear as a secondary phenomenon to preserve licensing context.

A striking effect of pied-piping is a V-to-V Raising in Belgian and Swiss dialects, see Haegeman (1992). In these dialects V-to-V Raising may pied-pipe its arguments. If one takes a verb in a VP_3 that is subcategorized for a particle, a prepositional phrase and a direct object, as [*de bal*_{DP} *uit de vijver*_{PP} *op*_{particle} *duiken*] ('the ball out (of) the pond up-dive'), a great variety of spurious distributions emerges. Raising the verb *duiken*₃ ('dive') to the causative *laten*₂ ('let') then yields the variants in (20).

- (20) a. *de bal uit de vijver op* – [*laten duiken*]
 b. *de bal uit de vijver* – [*laten opduiken*]
 c. *de bal* – [*laten uit de vijver opduiken*]
 d. – [*laten de bal uit de vijver opduiken*]
 (let the ball dive up out of the pond)

The verb *duiken*₂ ('dive') may pied-pipe or strand any of the three dependents. Haegeman (1992) analyses this by means of scrambling arguments out followed by VP-to-V raising. An alternative that avoids scrambling is given in Evers (2003: 83) by means of partial-VP Raising, seen as optional pied-piping. Pied-piping a subject of *opduiken* ('dive up'), say *haar kikker*^{3,4} ('her frog') in (2b), is absolutely out of the question. It is licensed only as a specifier of *laten opduiken* ('let dive up'), see (2b).

3.3 Long passives

Wurmbrand (2001) makes a distinction between functional verbs that restructure their infinitival complements and non-functional verbs that seem to do so. The latter {causatives, perceptuals, control verbs} assign a theta-role to a DP argument. Wurmbrand (2001: 125, fn.46) observes that all relevant clausal complements lack a tense/ I^0 marking. That implies, given the present-day theory, that such complements are transparent for argument movements and case-makings. By consequence, there is no need for a head-restructuring in order to derive argument-restructuring. Since the V^0 heads often seem to have been rearranged, this is best taken care of, according to Wurmbrand, by some set of ad hoc re-orderings in the phonological component. Head rearrangement in the syntactic component, she observes correctly, requires a general syntactic principle and this does not seem to be present. That is a correct and beautiful challenge.

Fortunately, Wurmbrand considers the West-Germanic verbal predicates as head-final. She also considers transitive predicates as inherently complex, such that the accusative is brought in as the external argument of a sub-predicate. From that point on, I reason as follows. Any external argument asks for the cooperation of two heads, the external head for the argument's case and the internal head for the argument's theta role. Both heads affect the external argument and a fusion of the heads as in V-to-I expresses this effectively. If no head-to-head fusion takes place (and other material appears between the two heads), the new external argument is still to the left

The V-cluster in (22) headed by the past participle will not be able to assign a structural case to the external argument of *hochzubringen*₃ ('to dive up') *den/der Frosch* ('the frog'), since by assumption the past participle will not assign a structural case. Nor can

the past participle V-cluster be expected to re-license or re-attach the accusative *den Ball*^{3,4} ('the ball'). It would have been better if the accusative *den Ball*^{3,4} had not been assigned by *aus*₄ – *hochzubringen*₃ ('to dive up') in view of the impending passive effect that will come in with the past participle *versucht*₂ ('tried'). The information of some passive effect should have been available in the cycle *aus*₄ – *hochzubringen*₃ such that the accusative assignment to *den Ball* could have been blocked, but it isn't. The cyclic distributions of grammatical information makes that impossible. The information of the passive effect comes in too late. Something like that holds as well for the past participle *versucht* itself. If the past participle is constructed with *haben* ('have') all passive past participle properties are cancelled.⁵ This information about the passive or active status of the past participle will also become available one cycle too late.

It could not have been that difficult for the West-Germanic Sprachgeist to save the system of local cyclic restructuring in (21b) as well as in (21a). The past participle passive *versucht*₂ ('tried') might have added an infinitival prefix to its complement *hochzubringen*₃ ('to dive up') in (21b) such that this complement could already adapt to the passive mode. The more serious question is whether the Sprachgeist is fundamentally committed to the system of cyclic local subject formation, and if so, how it handles the late information about the passive.

I propose the following way out. Rather than merging and moving our way up from right to left, we may as well move and merge our way down from left to right, like one expects of a parser. The general convention to start with the most embedded structure seems no more than a matter of exposition. One begins in classroom expositions with some lexical frame and then there are distributional variants for passive, wh-movement, V-second, Neg-placement and so on. Soon rules are busy to fill in the left periphery. This can be reversed. Rather than starting with a lexical theta-frame in order to find the appropriate structural case configurations, one may start with structural case configurations and derive how they find their way to lexically given frames. For example the structural cases in (21a) *der Frosch*^{1,2} and *den Ball*^{3,4} indicate subjects that refer to extended projections assembled in the complex predicate. The first structural case in (21a) *der Frosch*^{1,2} ('the frog') relates as a nominative to the extended projection [*versuchen*₂ *wird*₁] ('try will'). It gets the agentive relation defined by *versuchen*₂ ('try'). The control verb *versuchen*₂ *zu* ('try to') activates *zu* within the V-cluster of (21a). This *zu* can subsequently (i.e. after V-to-V Raising) be identified as

5. It makes sense that the past participle in Dutch does not come in as a restructuring predicative head at all. It is always and obligatorily turned in for the infinitive (the so-called *Infinitivus pro Participio*). German restricts this largely to an option for a restrictive class of restructuring verbs. The past participle passive is not used in Dutch V-clusters at all. See for an analysis Evers (2003: 77).

zu^{2,3}, but will require a matrix antecedent (the controller). The second structural case in (21a) *den Ball*^{3,4} ('the ball', accusative) relates to the projection extension [*aus*₄ *dem Teich hochzubringen*₃]₃. The left-hand part *aus*₄ gives the locality that determines the theta role of theme.

The passive variant in (21b) is much different, but fits the same mechanism. The structural case *der Ball* ('the ball', nominative) arrives at almost the same projection extensions, but [*versucht*₂ *werden*₁] fails to define a theta relation due to its past participle core at the left-hand side. The next extended projection [*hochzubringen*₃ *versucht*₂] ('to dive up tried') takes over. The left-hand side [*hochzubringen*₃] should determine the theta role. Now local information comes in about the relevance of *zu* in [*hochzubringen*₃]. The infinitival prefix *zu* has Haider's (1984) intransitivizing effect. [*hoch-zu-bringen*]₃ is not governed by an active verb, but by a past participle [*versucht*₂] and it intransitivizes, see (23).

- (23) a. *der Ball ist hoch*₄ *zu bringen*₃
 the ball is up to dive
 (the ball is to be dived up)
- b. *ein hoch*₄ *zu bringender*_{3,4} *Ball*
 an up to dive-agr ball
 (a to be dived up ball)

Because the infinitival prefix *zu* in (23) is not governed by an active verb, there is no access to the agent theta role of *hoch(zu)bringen*. The internal theta role of [*hoch*₄ *zu bringen*₃] is selected as the first theta role available. Had *hochzubringen*₃ in (23) been an intransitive verb, for example *treiben* ('float') or *schwimmen* ('swim') as in (24), the construction would have been ungrammatical for lack of a theta role. The intransitivizing *zu* in (24) removes the external theta role of *treiben*₃/*schwimmen*₃ and *der Ball* or *der Frosch* in (24) remains without interpretation from the successive partial predicates 2,3 and 1,2.

- (24) a. **der Ball wurde [zu treiben versucht]*
 the ball was [to float tried]
- b. **der Frosch wurde [zu schwimmen versucht]*
 the frog was to swim tried

The main point is that long passives may be derived by a cyclic local process of subject identification if the Sprachgeist is parser-oriented.

4. Epilogue

The survival time of a grammatical principle in generative grammar is not particularly high. After some twenty years it requires creative interpretation to show how some

idea still fits or no longer fits. The factual instability of linguistic concepts also indicates that claims about their psychological reality are not to be taken too seriously. Psychological reality is more a philosophical gesture. It is not meant to point to an existing ultimate testing ground for linguistic concepts. Psychological reality has the status of a Ding an Sich. Something that common sense dictates to be present, but also something that cannot be used to check cognitive distinctions beyond first constructing a tentative model of it. The strong point of mere grammatical analyses is rather the way they can mutually reinforce each other. They constitute a grammatical network and one cannot take elements out and ask for their independent existence somewhere else. The internal consistency of the network yields the feeling that it represents a reality that autonomously confirms or disconfirms our expectations. Further adaptations of the network may be partially revealing for psychological or neurological mechanisms, but they need not be. Structures and procedures taken up by the mind get a novelty of their own that cannot be understood by the initial means that are used to implement such structures or procedures. The piece of chalk used to get a mathematical proof on the blackboard does not contain axioms or a theory of proof. In the same vein, the hidden hand postulated by Adam Smith for market forces has been studied by economic models. One may claim that the economic models reconstruct the mind of those that visit the market, but that will further solve very little about the market problems themselves.

Things need not be much different with the notion psychological reality in grammar. Psychological reality may be traded in for Herder's Sprachgeist. Herder's notion seems more metaphysical, but the opposite is true. The notion Sprachgeist brings out that cultural artifacts, including grammar, take shape in human history and develop an internal coherence of their own that reveals more than the devices by which they are built up. That inner coherence is intuitively felt and can only be grasped tentatively. This is not to deny that the Sprachgeist may partly materialize in the human mind/brain. Such partial parallels are not necessary, though.

References

- Baker, M. 1988. *Incorporation. A Theory of Grammatical Function Changing*. Chicago IL: University of Chicago Press.
- Barbiers, S. 2005. Word order variation in three-verb clusters and the division of labour between generative linguistics and sociolinguistics. In *Syntax and Variation. Reconciling the Biological and the Social*, L. Cornips & K. Corrigan (eds), 233–264. Amsterdam: John Benjamins.
- Bech, G. 1955. *Studien über das Deutsche Verbum Infinitivum*. Kobenhavn: Ejnar Munksgaard.
- Burzio, L. 1986. *Italian Syntax*. Dordrecht: Kluwer.
- Chomsky, N. 1995. *The Minimalist Program*. Cambridge MA: The MIT Press.

- Evers, A. 2003. Verbal clusters and cluster creepers. In *Verb Constructions in German and Dutch*, P.A.M. Seuren & G. Kempen (eds), 43–89. Amsterdam: John Benjamins.
- Haegeman, L. 1992. *Theory and Description in Generative Syntax. A Case Study in West Flemish*. Cambridge: CUP.
- Haider, H. 1984. Was zu Haben Ist und Was zu Sein Hat - Bemerkungen zum Infinitiv. *Papiere zur Linguistik* 30: 23–36.
- Haider, H. 1993. *Deutsche Syntax, Generativ. Vorstudien zur Theorie einer Projektiven Grammatik*. Tübingen: Narr.
- Koopman, H. & Szabolcsi, A. 2000. *Verbal Complexes*. Cambridge MA: The MIT Press.
- Lieber, R. 1983. Argument linking and compounds in English. *Linguistic Inquiry* 14: 251–285.
- Mulder, R. 1992. The Aspectual Nature of Syntactic Complementation. Ph.D. dissertation, University of Leiden.
- Paardekooper, P.C. 1955. *Syntaxis, Spraakkunst en Taalkunde*. Den Bosch: Malmberg.
- Rutten, J. 1992. Infinitival Complements and Auxiliaries [Amsterdam Studies in Generative Grammar 4]. Ph.D. dissertation, University of Amsterdam.
- Ryle, G. 1947. *The Concept of Mind*. London: Hutchinson's University Library.
- Schmid, T. & Vogel, R. 2003. Dialectal variation in German 3-verb-clusters. A surface-oriented optimality theoretic account. *Journal of Comparative Germanic Linguistics* 7: 235–274.
- Seuren, P. 1996. *Semantic Syntax*. Oxford: Blackwell.
- Wurmbrand, S. 2001. *Infinitives: Restructuring and Clause Structure*. Berlin: Mouton de Gruyter.
- Zwart, C.J.W. 1996. *Morphosyntax of Verb Movement: A Minimalist Approach to the Syntax of Dutch*. Dordrecht: Kluwer.

Scope assignment in child language

On the role of the Question Under Discussion

Andrea Gualmini & Sharon Unsworth*
Utrecht University

This paper focuses on recent work on children's interpretation of scopally ambiguous sentences. We review current literature on the topic and we discuss two theoretical notions, namely the notion of *surface scope* and the notion of *Question Under Discussion* (QUD). We argue that both notions are theoretically motivated but pertain to different domains. In particular, in agreement with current literature, we acknowledge that notions defined over the syntax-semantics mapping are psychologically real and play a role in determining the number of interpretations that are available for any given sentence (see Fox 2000). However, we propose that inverse scope and surface scope are equally available to the psychological parser. In particular, we argue that ambiguity resolution is guided by contextual congruence rather than by considerations about the syntax-semantics mapping.

1. Scope resolution in child language

One of the most investigated topics in semantic theory is the interpretation and the interaction of scope-bearing elements (see Aoun & Li 1993; Fox 2000; May 1985 & Reinhart 1997). In recent years, scope ambiguities involving negation have also been the subject of many investigations of child language. In particular, much current research on this topic starts from the seminal work of Musolino (1998), which documented children's non-adult behavior with sentences such as the ones reported below.

- (1) Every horse didn't jump over the fence
- (2) The detective didn't find some guys
- (3) The detective didn't find two guys

*Many thanks to Christina Helder, Hannie van Rooijen, Arjanneke Bruggeman and Anneloes Canestrelli for data collection. Andrea Gualmini's contribution was supported in part by a Standard Research Grant from the Social Sciences & Humanities Research Council of Canada (SSHRC) and a VIDI grant from the Netherlands Organization for Scientific Research (NWO) and Utrecht University.

Each of the sentences above contains two operators: negation and a quantified noun phrase. This yields two logically possible scope assignments for each sentence. To illustrate, (1) is ambiguous between the two interpretations in (4) and (5).

- (4) Every horse is such that it did not jump over the fence
 $\forall x [\text{horse}(x) \rightarrow \neg \text{jump over the fence}(x)]$
- (5) Not every horse jumped over the fence
 $\neg \forall x [\text{horse}(x) \rightarrow \text{jump over the fence}(x)]$

The two interpretations of (1) listed in (4) and (5) result from the relative scope assignment to negation and *every*, as suggested by the order of the logical operators \neg and \forall in the logical formulae. In the semantic literature, the interpretation in (4) is referred to as the *surface scope* interpretation of (1), while (5) is referred to as the *inverse scope* interpretation of (1). This is because in the interpretation in (4), the scope bearing elements *every* and *not* are interpreted in the same order with which they appear in the overt syntax, whereas in (5) they are interpreted in the opposite order.

The research question that Musolino (1998) and many others following him have addressed is whether the two interpretations of a scopally ambiguous sentence have equal status in children's grammars. More generally, Musolino (1998) wanted to find out whether young children are capable of accessing both the surface scope and the inverse scope interpretation of sentences containing negation and another scope-bearing element.

To address this question, Musolino (1998) conducted a series of experiments employing the Truth Value Judgment task that tested children's interpretation of sentences like (1)–(3). The experimental evidence collected by Musolino (1998) suggests that inverse and surface scope interpretations of negative sentences are not equally accessible to young (4- and 5-year-old) children. In fact, the experimental evidence collected by Musolino (1998) suggests that for the sentences in (1)–(3), children resort to their surface scope interpretations, paraphrased below.

- (6) Every horse is such that it did not jump over the fence
 $\forall x [\text{horse}(x) \rightarrow \neg \text{jump over the fence}(x)]$
- (7) It is not the case that the detective found some guys
 $\neg \exists x [\text{guy}(x) \rightarrow \wedge \text{detective found}(x)]$
- (8) It is not the case that the detective found two guys
 $\neg \exists x [\text{two guys}(x) \rightarrow \wedge \text{detective found}(x)]$

To review the experimental findings briefly, Musolino (1998) showed that 4- and 5-year-old children would consistently reject a sentence like (1) in a context in which two horses had jumped over the fence and one horse had stayed behind. Furthermore, Musolino (1998) found that children as old as 5;9 rejected sentences like (2) or (3) if

the detective found two of the four guys available in the context, whereas a group of adult controls consistently accepted the target sentences. More precisely, many of the children tested by Musolino (1998) claimed that (2) or (3) was incorrect because the detective had indeed found some guys.

Children's failure to access the inverse scope interpretation of the relevant sentence, even though that interpretation would have made the sentence true, was taken as evidence that children's semantic scope is limited to overt syntactic scope. This is the Observation of Isomorphism proposed by Musolino (1998).

As extensively discussed by Musolino (1998), Musolino (2006) and Musolino, Crain and Thornton (2000), children's non-adult behavior as described above can be approached in two different ways. As a first possibility, children might not have access to the same range of interpretations as adults. On this scenario, children's non-adult behavior might ultimately be due to a problem with any of the factors that determine what interpretations are available to adults. For instance, as we will review momentarily, Fox (2000) has observed that when inverse scope interpretations are indistinguishable from surface scope interpretations, only the latter option is available. Thus, surface scope plays a crucial role, in that it provides a baseline against which inverse scope interpretations are evaluated. Given Fox's account, one might put forward the hypothesis that inverse scope is subject to even more stringent grammatical constraints in child language. As a second possibility, children and adults might have access to the same interpretations, but they might differ in which interpretation they take to be the speaker's intended interpretation.

As it turns out, the evidence documented by subsequent research speaks in favor of the second scenario. In particular, recent studies have shown that children's interpretation of sentences like (1)–(3) is affected by the context. For present purposes, an interesting pattern emerges once we consider a study by Gualmini (2004). This investigation included stories in which a character had a task to carry out. In one of the trials, children were told a story about a troll who was supposed to deliver four pizzas to Grover. On the way to Grover's house, two pizzas fell off the truck, and the troll only managed to deliver two of the four pizzas. Children were then asked to judge either (9) or (10).

- (9) The troll didn't deliver some pizzas
- (10) The troll didn't lose some pizzas

The results documented in Gualmini (2004) show that children respond differently to the two sentences. Thirty 4- and 5-year-olds participated in the experiment. Children accepted the inverse scope interpretation of (9) to a higher extent than they did for (10) (i.e. 90% and 50% respectively).

A further study by Gualmini, Hacquard, Hulsey and Fox (2005) showed that, just as in the case of (9), the same context used by Gualmini (2004) also leads children to

access the inverse scope interpretation of sentences equivalent to (1) and (3), namely (11) and (12), respectively (for data on sentences containing the universal quantifier *every*, see also Musolino & Lidz 2006).

- (11) Every pizza wasn't delivered
- (12) The troll didn't deliver two pizzas

In other words, the experimental manipulation discovered by Gualmini (2004) leads children to access inverse scope interpretations for all the scope-bearing elements that had yielded non-adult behavior in Musolino's (1998) investigation.

Let us sum up. We started with children's non-adult behavior as documented by Musolino (1998). Following previous work, we noted that children's behavior could follow from a problem in the generation of the adult interpretation(s) or from a problem in the selection of that interpretation (out of the available options). We then noted that the data suggest that children's grammars generate all the relevant interpretations, and the source of the problem seems to lie in how children go about selecting the intended interpretation. The question we would like to address is whether this change in the phenomenon under investigation also calls for a change in the theoretical machinery that is needed to explain the facts. Now that the problem seems to lie in the scope interpretation which children choose (rather than the ones they can generate), we need to consider whether theoretical notions other than the notion of surface scope are needed to explain the facts.

2. On the relevance of surface scope

In this section, we focus on one notion that has attracted much attention in the debate on scope resolution in child language, namely the notion of surface scope. In agreement with much current research in theoretical linguistics, we acknowledge the importance of this notion. In particular, the surface scope interpretation is used in explaining why logically plausible interpretations are often unavailable (see Fox (2000)). However, contra much child language literature, we also argue that the notion of surface scope does not explain how children go about selecting the intended interpretation (out of the alternatives that are available to them).

Let us first consider one phenomenon for which surface scope plays a crucial role. Compare the following examples, discussed in Fox (2000):

- (13) A girl admires every teacher
- (14) Every boy admires every teacher

If we consider (13), one can think of truth-conditional evidence for the existence of the inverse scope interpretation. In particular, one can think of situations in which only

the inverse scope interpretation holds. For instance, sentence (13) is true in a situation where Mr. Sumner is admired by Jenny, Mrs. Townsend by Zoe and Miss Linklater by Lucy (and there are no other girls who admire the teachers). In such a situation, only the inverse scope interpretation of (13) is true. Due to the entailment relations that hold between the two readings, the issue remains whether the surface scope interpretation is also available, an issue that will be resolved shortly, but at the very least one knows that the inverse scope interpretation needs to be posited. When it comes to (14), however, the surface and inverse scope interpretations happen to be true in exactly the same range of situations, that is, where Marc admires Mr. Sumner, Mrs. Townsend and Miss Linklater, and so do Robert and John. In other words, the two putative scope assignments of (14) are truth-conditionally indistinguishable. The consequence is that we do not know which scope assignments are possible.

To make up for the shortage of truth-conditional evidence for the existence of both scope assignments, Fox (2000) shows that one can turn to other tests. For instance, Fox (2000) makes use of a constraint on ellipsis. Consider (15).

- (15) A girl admires every teacher. Every boy does, too

What is interesting about (15) is that the antecedent sentence, which is identical to (13), can only receive a surface scope interpretation. Thus, the sentence now seems to require that there be a single student, say Jenny, who admires all the teachers, that is, Mr. Sumner, Mrs. Townsend and Miss Linklater. The question is why the inverse scope interpretation, which is available for (13), disappears when the same sentence is used in the antecedent sentence of an ellipsis construction. In other words, we need to explain why whereas (13) is true even in a situation in which every teacher is admired by a different girl, for (15) to be true, it must be the case that the very same girl admires every teacher (and, of course, that every boy admires every teacher, too).

According to Fox (2000), an explanation emerges once we notice that ellipsis constructions are subject to a Parallelism constraint: the elided verb phrase needs to be identical to a phonologically-realized verb phrase in the antecedent sentence. A consequence of Parallelism is that the scopal relationship of the elided sentence must be identical to the one of the antecedent sentence. In short, the antecedent and the elided sentence must receive parallel scope assignment. On this view, the explanation for the interpretation of (15) is quite simple: knowing that the antecedent sentence and the elided sentence must receive parallel scope, and seeing that in this particular case the antecedent sentence can only receive a surface scope interpretation, it must be the case that the sentence containing ellipsis must be interpreted on its surface scope interpretation. In a sense, the antecedent sentence takes the only scope assignment which is available for the elided sentence. For present purposes, however, what is crucial is not the fact that the antecedent sentence must have the same scope assignment as the elided sentence, but rather that this provides us with a window into the scope assignment of

the elided sentence itself. And what we see is that only the surface scope interpretation is available.

The consequences of Fox's (2000) investigation are far reaching. The generalization proposed by Fox (2000) is that whenever the two logically possible interpretations of a sentence are truth-conditionally indistinguishable, only surface scope needs to be posited. In turn, this means that inverse scope interpretations – but not surface scope interpretations – are subject to a constraint: they must yield an interpretation that is truth-conditionally distinguishable from the surface scope interpretations. This is a generalization that Fox (2000) attempts to derive from the non-economical nature of scope-shifting operations (see also Reinhart (2006)).

Let us take stock. The relevance of the notion of surface scope is indisputable. The preceding discussion highlights one constraint on the generation of inverse scope interpretations. At first glance, one could hypothesize that Musolino's original findings follow from children's grammar encoding a yet stronger constraint than the one discussed by Fox (2000), maybe even a total ban against inverse scope interpretations. Under this scenario, surface scope would define the hypothesis space available to children. Nevertheless, the results documented by Gualmini (2004) and Gualmini, Hacquard, Hulsey and Fox (2005) falsify this hypothesis. We now know that children's grammars can generate both scope interpretations of sentences like (1)–(3).¹ The next question is whether the notion of surface scope can be invoked in any other way to explain children's non-adult behavior. This is the strategy followed by Musolino and Lidz (2006), who argue that a parsing preference for surface scope, together with a preference for true interpretations, accounts for children's non-adult behavior. We, however, would like to pursue a different strategy. In fact, there are at least two pieces of evidence suggesting that the notion of surface scope is not useful in approaching this last phenomenon.

The first piece of evidence against the role of surface scope in children comes from Dutch-speaking children. Consider the Dutch example in (16), taken from Krämer (2000).

- (16) De jongen heeft een vis niet gevangen
 The boy has a fish not caught
 'There is a fish the boy hasn't caught.'

Children's interpretation of sentences like (16) was investigated by Krämer (2000). The experimental results show that 38 Dutch-speaking children from 4;0 to 7;7 rejected

1. In fact, some data we review shortly suggest that children can even select inverse scope interpretations more freely than adults (see also Sano 2004 and Yamakoshi & Sano 2007).

(16) as a description of a story in which a boy had caught two fish out of the three fish available in the context, whereas adults always accepted it (see also Unsworth (2005)). Children, unlike adults, apparently interpreted the indefinite *een vis* (*one fish*) in the scope of negation, which corresponds to the inverse scope interpretation of (16). Even if we simply take Isomorphism as a descriptive generalization about how differences between children and adults can manifest themselves, the challenge posed by the Dutch data is quite clear: unlike many of the previous findings with English-speaking children, Dutch-speaking children select the *inverse scope* interpretation more readily than the surface scope interpretation.

An additional piece of evidence against Isomorphism comes from English-speaking children and is reported in a study by Hulsey, Hacquard, Fox and Gualmini (2004). These authors designed an experiment to tease apart the relative contribution of a putative preference for surface scope and contextual information. We will return to their particular model of how context guides ambiguity resolution momentarily. For now, we consider the relevance of their experiment for the role of surface scope in ambiguity resolution. Basically, Hulsey et al. reasoned as follows. If a preference for surface scope plays a crucial role in explaining children's non-adult behavior, a manipulation of the surface syntax of the stimuli should have an effect on children's behavior. In an experiment, Hulsey et al. (2004) tested children's interpretation of sentences such as (17) and (18).

(17) Some pizzas were not delivered

(18) Some pizzas were not lost

If we compare sentences (17) and (18) with the sentences used by Gualmini (2004) (i.e. (9) *The troll didn't deliver some pizzas* and (10) *The troll didn't lose some pizzas*), we see a different surface syntax, namely passives versus actives, respectively. Thus, if surface syntax dictates children's responses, children should respond to (17) differently from the way they respond to (9) and they should respond differently to (18) and (10). This prediction turned out to be incorrect, however. Hulsey et al. discovered that children in the relevant developmental stage interpret (17) like (9) and (18) like (10). In other words, children's responses are not dictated by structural form.

To illustrate, let us focus on children's interpretation of sentences like (18). The finding documented by Hulsey et al. (2004) was that adult speakers of English always accepted sentences such as (18) in the context under investigation, thereby consistently selecting its true surface scope interpretation (i.e. there are some pizzas that were not lost). By contrast, children rejected the target sentence about half of the time. In particular, about half of the time, children's responses were indicative of the inverse scope interpretation (i.e. no pizza was lost). Thus, some of Hulsey's et al findings resemble Krämer's (2000) and Unsworth's (2005) findings, in that children select the inverse scope interpretation for sentences that adults seemingly must interpret according to

surface scope. In our view, this undermines the hypothesis that surface scope plays a crucial role in determining which interpretation of a scopally ambiguous sentence children select.

To sum up, in this section we have discussed the notion of surface scope. We do not dispute the importance of such a notion in explaining which interpretations are generated in the grammar of children or adults. However, in our view, the data suggest that the notion of surface scope is of little use in explaining which interpretation is ultimately selected by children's (or adults') parser. To answer this question, we would like to invoke a different notion.

3. On the relevance of the Question Under Discussion

We concluded the previous section by arguing that the notion of surface scope, though theoretically useful, does not help us in solving the problem outlined in Section 1, namely why children and adults sometimes select a different interpretation of scopally ambiguous sentences. In this section, we turn to an independently motivated notion that we believe can be used to answer this question. This is the Question Under Discussion (QUD).

The relevant study is due to Hulsey, Hacquard, Fox and Gualmini (2004). These authors developed a new model of scope resolution in child language that makes reference to independently motivated principles of communication. According to this model, which Hulsey et al. (2004) call the Question-Answer Requirement (QAR), children select the scope assignment that allows them to address the Question Under Discussion. In turn, according to Hulsey et al. (2004), a good answer to a question is a proposition that entails an answer to that question.

According to Hulsey et al. (2004), what is relevant in the pizza story used by Gualmini (2004) is the troll's task. At the end of the story, one wants to know whether the troll has carried out his task or not. This amounts to asking the 'yes/no' question *Did the troll deliver all the pizzas?*. Let us recall the relevant target sentences.

- (9) The troll didn't deliver some pizzas
- (10) The troll didn't lose some pizzas

Notice that both inverse and surface scope interpretations of (9) entail an answer to that question. On the surface scope interpretation, the troll didn't deliver any pizzas, hence the answer to the QUD is negative. On the inverse scope interpretation, there are some pizzas that the troll didn't deliver and therefore the answer to the QUD is also negative. Thus, as far as the QAR is concerned, either scope assignment is viable, and children can make use of the Maxim of Charity (see Grice (1975)), which essentially means that they will select the interpretation that makes the target sentence true (i.e. inverse scope). By contrast, in the case of (10), only the surface scope interpretation addresses the

contextually relevant question. In a context in which delivering and losing pizzas are the relevant alternatives, the surface scope interpretation of (10) is equivalent to the proposition paraphrased in (19), which is false in the context, but answers the QUD.

- (19) The troll delivered all the pizzas

By contrast, the inverse scope interpretation of (10) in the present context is equivalent to (20), which implicates an answer to the QUD, but does not entail it. In particular, (20) is true both if the troll delivered all of the pizzas and if he only delivered some of them. Thus, by simply knowing that (20) is true, one does not know the answer to the question whether the troll delivered all the pizzas.

- (20) The troll delivered some pizzas

Given that only one interpretation satisfies the QAR, the Maxim of Charity ends up being violated. In other words, Hulsey et al. (2004) argue, if only one interpretation addresses the QUD, that interpretation is selected regardless of whether it makes the target sentence true or false. Moreover, whether an interpretation corresponds to surface or inverse scope of the target sentence does not play a role in determining whether that interpretation is selected.

A prediction of this model is that, for any given context and for any given predicate, children will prefer the same interpretation, regardless of whether it amounts to surface scope or inverse scope. As we saw above, this prediction was corroborated by Hulsey et al. (2004) in an experiment testing children's interpretation of the following sentences:

- (21) Some pizzas were not delivered
(22) Some pizzas were not lost

The crucial case is given in (22). As we saw above, if surface scope figures prominently in scope resolution, children should consistently interpret (22) on its true surface scope interpretation. The predictions of QAR are different, however. This model does not assign any privileged status to surface scope in ambiguity resolution. When a sentence is ambiguous, the listener must consider whether any interpretation addresses the Question Under Discussion. What is important is the proposition expressed by each reading of an ambiguous sentence, not whether or not that interpretation differs from surface syntax. Within the present context, the Question Under Discussion is whether the troll delivered all the pizzas. As in the case of (10), the only interpretation of (22) that is a good answer to that question is the interpretation in which negation takes scope over *some*. In the active sentence, this is the surface scope interpretation, but in a passive sentence such as (22), this is the inverse scope interpretation. Therefore, the QAR predicts that some children will reject (22). More precisely, the QAR predicts that children for whom the sentence is ambiguous will reject (22) to the same extent to which they rejected (10) and for the same reason that they rejected (10). In this

context, some children should resort to the interpretation in which negation has scope over *some* because only that interpretation addresses the Question Under Discussion.

The results show that English speaking 3- and 4-year-old children always accepted (21), but many of them rejected (22). In particular, half of the subjects tested by Hulsey et al. (2004) rejected (22) on the grounds that some pizzas were indeed lost, thereby accessing the inverse scope interpretation of (22) (i.e. it is not the case that some pizzas were lost). As predicted by the QAR model, the rate of rejection for (22) (i.e. 57%) closely mirrors the rate of rejection found by Gualmini (2004) for (10) (i.e. 50%).

An additional prediction of the QAR account proposed by Hulsey et al. (2004) pertains to other constructions investigated by Musolino (1998). Having seen how the QAR account can explain children's documented non-adult behavior for sentences containing the indefinite *some* in object position, we can ask whether the QAR account can also explain children's non-adult behavior for sentences containing the indefinite *two* in object position or the universal quantifier *every* in subject position. As we saw above, this question was addressed in a later study by Gualmini et al. (2005). In particular, these authors showed that the same contextual manipulation discovered by Gualmini (2004) leads children to access the inverse scope interpretation of sentences containing negation and the indefinite object *two* (e.g. *The troll didn't deliver two pizzas*) as well as sentences containing the universal quantifier *every* and negation (e.g. *Every pizza wasn't delivered*) to a higher extent than observed in previous literature (see Gualmini (2005/2006) for a summary of the data and Gualmini (2007) for further discussion about how the QAR can explain previous data by Musolino (1998)).

The QAR account also invites a fresh look at the putative cross-linguistic difference that emerges once we compare English- and Dutch-speaking children. As we wrote in Section 2, Dutch-speaking children's behavior provided us with the first piece of counterevidence to the Observation of Isomorphism, in that English- and Dutch-speaking children's non-adult behavior seemed to surface in such different ways. Recall the Dutch example in (16) repeated below as (23).

- (23) De jongen heeft een vis niet gevangen
 The boy has a fish not caught
 'There is a fish the boy hasn't caught.'

As we saw above, Krämer (2000) found that children as old as seven interpreted (23) on its inverse scope interpretation, whereas adults consistently resorted to surface scope. This gives the illusion of a puzzle, as Dutch- and English-speaking children's non-adult behavior seems to take different shapes. The puzzle disappears, however, once we abandon Isomorphism and we consider the predictions of the QAR. One possibility is that the question conveyed by the story in Krämer's original experiment is something like *Did the boy catch some fish?*. This question would only be addressed by the inverse

scope interpretation of the indefinite, i.e. it is not the case that the boy caught a fish. The surface scope interpretation, namely that there is a fish which the boy did not catch, does not address this question because knowing that there is a fish which the boy did not catch does not tell us whether there are any fish which he *did* catch. The prediction of the QAR is that the acceptance of the target sentence should increase if the context makes prominent the following question: *Did the boy catch all the fish?* As it turns out, this prediction was borne out.

In a series of experiments, Unsworth and Helder (to appear) and Unsworth and Gualmini (2007) show that making the same experimental manipulation discovered by Gualmini (2004) leads Dutch-speaking children to access the surface scope interpretation for sentences such as (23) more readily than observed in previous research. In Krämer's original experiment, the experimental scenario introduced a boy and some fish which the experimenter suggested the boy was going to catch, which, as noted above, might lead to the QUD *Did the boy catch some fish?*. Unsworth and Helder (to appear) and Unsworth and Gualmini (2007) modified this scenario such that the boy had a particular task to carry out, that is, he had to catch all the fish in order to win a fishing competition. As a result, the Question Under Discussion was "*Did the boy catch all the fish?*". Both the surface and inverse scope interpretations of the sentence in (23) answer this question.² On the surface scope interpretation, there is a fish which the boy did not catch and so Question Under Discussion receives an answer. Similarly, on the inverse scope interpretation the Question Under Discussion is answered because if it is not the case that the boy caught a fish, then it is also not the case that he caught all the fish. Assuming that children follow the Maxim of Charity, they should select the interpretation which makes the test sentence true, that is, the surface scope interpretation. This is exactly what they did.

Unsworth and Gualmini (2007) tested monolingual Dutch-speaking 5-year-olds using a Truth Value Judgment task using stories (five in each condition) in which the main character had to carry out a task involving all of the objects of a given set, such as catching all the fish as part of a fishing competition. In the first condition, the boy caught two out of the three available fish. As indicated above, in this condition both the

2. Following Krämer (2000: 106), the sentences in (23) and (24) were pronounced with the most natural intonation for the expected adult interpretation, that is, with the NP destressed and a slight stress on negation, and the indefinite article was pronounced in its unstressed form, i.e. *een* instead of *één*. Whilst it is true that in production, native speakers tend to produce scrambled indefinites as *één* (Unsworth 2005), the results for the native controls here demonstrate that this difference does not affect comprehension: scrambled indefinite objects are consistently interpreted as taking wide scope over adverbials such as negation when pronounced in their unstressed form.

surface scope and inverse scope interpretations answer the QUD. Seventeen children (5;0 – 6;0, $M = 5;6$) were tested in this condition. They accepted the surface scope interpretation 71% (57/80) of the time. This contrasts starkly with the 23% acceptance rate which Krämer (2000) observed for similarly-aged children.

In a second condition, similar to Gualmini (2004) and Hulsey et al. (2004), the boy caught just one of the available fish and fourteen children (5;4 – 5;11, $M = 5;7$) were presented with sentences such as (24).

- (24) De jongen heeft een vis niet laten zwemmen
 The boy has a fish not leave swim
 ‘There is a fish the boy didn’t leave.’

The surface scope interpretation of this sentence fails to answer the QUD: if we know that there is a fish which the boy has left swimming in the pond, we do not know whether he has caught all the fish. It is only the inverse scope interpretation which answers the QUD here: if we know that it is not the case that the boy left a fish, then, based on the contextual alternatives, we also know that he must have caught them all. The QAR thus predicts that children will accept the surface scope interpretation in this condition much less frequently than in the first condition. This prediction was borne out: the surface scope prediction was accepted for sentences such as (24) just 45% (38/85) of the time. The difference between the two conditions is statistically significant ($t = 2.032$, $p = .051$). The acceptance rate for adult controls was 100% (50/50) and 95% (38/40) in the two conditions, respectively.

In a similar experiment, Unsworth and Helder (to appear) replicated an experiment originally carried out with English-speaking children by Miller and Schmitt (2004). By employing experimental changes similar to Gualmini (2004), Miller and Schmitt (2004) showed that, in contrast to earlier results such as those discussed in Section 1, English-speaking 4-year-old children consistently selected the inverse scope interpretation of sentences like (25) to the same extent as adults.

- (25) Timothy didn’t blow out a candle

These authors attribute this finding to their use of objects belonging to a pre-defined set, e.g. the candles on a birthday cake, the drawers in a chest, etc. in the experimental scenarios. It is possible, however, that a second change which these authors implemented, namely requiring the protagonist in the story to carry out an action on all the objects (e.g. candles, drawers, etc.), may also have contributed to this result. Interpreted from the QAR approach, the fact that the protagonist had a particular action to carry out highlights the relevant QUD, that is, whether the task (e.g. blowing out all the candles, closing all the drawers, etc.) was carried out successfully. In doing so, this may warrant children’s selection of the inverse scope interpretation as this is the only one which provides a true answer to the QUD.

Unsworth and Helder (to appear) demonstrate that implementing the same experimental changes with Dutch-speaking children achieves a similar result. A Dutch version of Miller and Schmitt's (2004) original experiment was presented to two groups of children: 4-year-olds ($n = 15$; 4;2 – 4;11, $M = 4;6$) and 6-year-olds ($n = 15$; 6;5 – 7;0, $M = 6;8$). In each of the four trials, children were presented with stories in which a protagonist had to carry out an action on all members of a group of objects belonging to a pre-defined set (cf. Krämer's (2000) original experiment, replicated in Unsworth (2005), where this was not the case). Both groups of children were found to readily accept the inverse scope interpretation of sentences such as (23): 4-year-olds at a rate of 76.8% (46/60) and 6-year-olds all the time (80/80).

Taken together, the results from Unsworth and Helder (to appear) and Unsworth and Gualmini (2007) suggest that although Dutch- and English-speaking children may differ in how quickly they acquire (and unlearn) an ambiguity, once we look at how scope ambiguities are resolved, Dutch- and English-speaking children behave in the same way: they select the contextually relevant interpretation, regardless of whether this amounts to surface scope or inverse scope.

We would like to conclude by commenting on the notion of Question Under Discussion (QUD), which figures prominently in the QAR account by Hulsey et al. (2004). In particular, we would like to stress that although such a notion did not enter the debate on scope resolution until recently, the Question Under Discussion has been invoked to explain many other phenomena. Among others, the QUD seems indispensable in theories of focus and theories of implicatures (see van Rooij (2003) and Roberts (2004) among others for theoretical work).

A recent psycholinguistic investigation that makes use of the QUD was carried out by Zondervan (2007). This study contains the results of two experiments investigating the role of the Question Under Discussion for adults' computation of scalar implicatures (see Grice (1975) and Horn (1989)). The importance of contextual information in the computation of scalar implicatures is a widely discussed phenomenon, most often invoked in the Relevance Theory literature (see Carston (1998)). Zondervan (2007) focused on a different issue, namely whether the notion of QUD can be used to model the role of contextual information for the computation of scalar implicatures. Building on work by van Kuppevelt (1996) and van Rooij (2002), Zondervan (2007) set out to test the hypothesis that implicatures arise if and only if the scalar term is in a constituent that answers the QUD. In one of the experiments, the QUD was presented overtly. To illustrate, subjects heard the target sentence (e.g. *Harry brought bread or chips*), which was offered as a response to two different kinds of questions as illustrated below.

- (26) A: Who brought bread or chips?³
 B: Harry brought bread or chips.

3. The experiment was conducted in Dutch with adult speakers of Dutch.

- (27) A: What did Harry bring?
B: Harry brought bread or chips.

In both experimental conditions, the final outcome included Harry bringing both bread and chips. Thus, if subjects calculate the relevant implicature (i.e. Harry brought bread or chips *but not both*), the final outcome would warrant the rejection of the target sentence. However, the results show a significant difference in the calculation of implicatures: 55% for (26) and 73% for (27). Thus, when the scalar item *or* occurs in the constituent that is being questioned, subjects are more likely to compute the implicature and reject the target sentence than when the same scalar item occurs outside of the constituent being questioned. Interestingly, the same pattern emerges in a second experiment, in which the QUD is not presented overtly but rather needs to be evinced by the context. This suggests that the notion of Question Under Discussion can be independently motivated, thereby making it possible to ask whether that notion can be useful in explaining other phenomena. As should now be clear, we believe that in the case of scope resolution, the answer to this question is affirmative.

An anonymous reviewer invites us to comment on the relationship between the Question-Answer requirement (QAR) and the Question Under Discussion (QUD). As this paper demonstrates, both theoretical and psycholinguistic evidence can be used to underpin the notion of Question Under Discussion. Nevertheless, it is not the case that this notion alone explains children's behavior with scopally ambiguous sentences. In particular, assuming that every sentence which children (and adults) hear needs to be evaluated against the contextually relevant QUD, the next issue which needs to be addressed is which criteria guide such an evaluation. One hypothesis is that the parser filters the relevant interpretations by appealing to the notion of entailment: the parser privileges the interpretations that entail an answer to the QUD. This is the claim of the QAR model put forward in Hulsey et al. (2004). This is not the only possibility, however. In principle, it is possible that when it comes to other phenomena, the parser adjudicates among alternatives interpretations making use of weaker notions, such as implicature. Indeed, Hulsey et al. (2004) speculate that the difference between children and adults might ultimately boil down to them adopting different criteria to determine what counts as a good answer to the Question Under Discussion. On this scenario, children first and foremost select the interpretation that entails an answer to the Question Under Discussion, whereas adults might even make use of interpretations that implicate an answer such question (but see Zondervan, Meroni and Gualmini (2008) for discussion). In other words, the QUD and QAR are two distinct notions: the QAR puts forward a hypothesis about how the QUD can be used in ambiguity resolution.

To conclude, in this section we reviewed previous studies that make use of the notion of Question Under Discussion (QUD). We have started from the question of how scopal ambiguities are resolved and we have showed how the QUD can be used in explaining the findings documented in the literature on child language and motivated

further experimentation. Furthermore, we illustrated how the very same notion can be useful in explaining a seemingly unrelated phenomenon, namely the calculation of scalar implicatures in adults.

4. Conclusion

In this paper we have discussed two theoretical constructs: surface scope and Question Under Discussion. Although we agree with much previous literature on the importance of surface scope, we also showed how its relevance needs to be circumscribed. In particular, surface scope interpretations provide the baseline against which logically possible interpretations are evaluated. For instance, in the case of a sentence like (13) (i.e. *A student admires every girl*) the inverse scope interpretation is available, because it is semantically distinguishable from the more economical surface scope interpretation. However, as far as ambiguity resolution goes, whether or not a given interpretation is a surface scope interpretation is irrelevant. Thus, once we focus on the task of selecting which interpretation of (13) could be intended in a given context, both interpretations are on equal footing. The so-called computational system might care about the mapping between syntax and semantics, but the parser apparently does not. Within the domain of scope resolution, the parser seems to be guided by other considerations, in particular by the truth and – above all – by the discourse congruence of each interpretation.

References

- Aoun, J. & Li, Y.-H.A. 1993. *Syntax of Scope*. Cambridge MA: The MIT Press.
- Carston, R. 1998. Informativeness, relevance and scalar implicature. In *Relevance Theory: Applications and Implications* [Pragmatics & Beyond New Series 37], R. Carston & S. Uchida (eds), 179–236. Amsterdam: John Benjamins.
- Fox, D. 2000. *Economy and Semantic Interpretation*. Cambridge MA: The MIT Press.
- Grice, P. 1975. Logic and conversation. In *Syntax and Semantics*, P. Cole & J. Morgan (eds), 41–58. New York NY: Academic Press.
- Gualmini, A. 2004. *The Ups and Downs of Child Language: Experimental Studies in Children's Knowledge of Entailment Relationships and Polarity Phenomena*. New York NY: Routledge.
- Gualmini, A. 2005/2006. Some facts about quantification and negation one simply cannot deny: A reply to Gennari & Macdonald (2005/2006). *Language Acquisition* 13: 363–370.
- Gualmini, A. 2007. Scope resolution and overt questions: A test for the QAR. In *Proceedings of the Eighth Tokyo Conference on Psycholinguistics*, Y. Otsu (ed.), 121–135, Tokyo: Hituzi Syobo Publishing.
- Gualmini, A., Hacquard, V., Hulsey, S. & Fox, D. 2005. Beyond the surface of child language. Paper presented at the *Annual Meeting of the Linguistic Society of America*. San Francisco CA.
- Horn, L.R. 1989. *A Natural History of Negation*. Chicago IL: University of Chicago Press.

- Hulsey, S., Hacquard, V., Fox, D. & Gualmini, A. 2004. The question-answer requirement and scope assignment. In *Plato's problem: Problems in Language Acquisition*, A. Csirmaz, A. Gualmini & A. Nevins (eds), 71–90. Cambridge MA: MITWPL.
- Krämer, I. 2000. Interpreting Indefinites. Ph.D. dissertation, Utrecht University.
- Kuppevelt, J. van. 1996. Inferring from topics. Scalar implicatures as topic-dependent inferences. *Linguistics and Philosophy* 19: 393–443.
- May, R. 1985. *Logical Form: Its Structure and Derivation*. Cambridge MA: The MIT Press.
- Miller, K. & Schmitt, C. 2004. Wide-scope indefinites in English child language. In *Proceedings of GALA*, J. van Kampen & S. Baauw, 317–328. Utrecht: LOT Occasional Series.
- Musolino, J. 1998. Universal Grammar and the Acquisition of Semantic Knowledge: An Experimental Investigation into the Acquisition of Quantifier-Negation Interaction in English. Ph.D. dissertation, University of Maryland.
- Musolino, J. 2006. Structure and meaning in the acquisition of scope. In *Semantics in Acquisition*, V. van Geenhoven (ed.), 141–166. New York NY: Springer.
- Musolino, J., Crain, S. & Thornton, R. 2000. Navigating negative quantificational space. *Linguistics* 38: 1–32.
- Musolino, J. & Lidz, J. 2006. Why children aren't universally successful with quantification. *Linguistics* 44: 817–852.
- Reinhart, T. 1997. Quantifier scope: How labor is divided between QR and choice functions. *Linguistics and Philosophy* 20: 335–397.
- Reinhart, T. 2006. *Interface Strategies: Reference-set Computation*. Cambridge MA: The MIT Press.
- Roberts, C. 2004. Context in dynamic interpretation. In *Handbook of Contemporary Pragmatic Theory*, L.R. Horn & G. Ward (eds), 197–220. Oxford: Blackwell.
- Rooij, R. van. 2002. Relevance implicatures. Ms, ILLC, Amsterdam. Online: <<http://semanticsarchive.net/Archive/WIyOWUyO/Implicfinal.pdf>>.
- Rooij, R. van. 2003. Questioning to resolve decision problems. *Linguistics and Philosophy* 26: 727–763.
- Sano, T. 2004. Scope relations of QPs and scrambling in the acquisition of Japanese. In *Proceedings of GALA*, J. van Kampen & S. Baauw, 421–431. Utrecht: LOT Occasional Series.
- Unsworth, S. 2005. *Child L2, Adult L2, Child L1: Differences and Similarities. A Study on the Acquisition of Direct Object Scrambling in Dutch*. Utrecht: LOT Dissertation Series.
- Unsworth, S. & Gualmini, A. 2007. Uncovering the pattern of children's interpretation of negation and indefinites. Paper presented at BUCLD 32, Boston University.
- Unsworth, S. & Helder, C. To appear. Dissolving a Dutch delay: The case of specific indefinites. In *Proceedings of Generative Approaches to Language Acquisition 07*, M. João Freitas & A. Gavarró (eds). Cambridge: Cambridge Scholars Publishing.
- Yamakoshi, K. & Sano, T. 2007. Children's interpretations of some/every interaction in mono-clausal and bi-clausal structures in Japanese. In *Proceedings of the 2nd Conference on Generative Approaches to Language Acquisition North America (GALANA)*, A. Belykova, L. Meroni & M. Umeda (eds), 471–482, Somerville MA: Cascadilla.
- Zondervan, A. 2007. Effects of question under discussion and focus on scalar implicatures. In *Proceedings of the Fifth Semantics in the Netherlands Day*, M. Kluck & E.J. Smits (eds), 39–52. Groningen: CLCG.
- Zondervan, A., Meroni, L. & Gualmini, A. 2008. Experiments on the role of the question under discussion for ambiguity resolution and implicature computation in adults. In *Proceedings of SALT 18*, T. Friedman & S. Ito (eds), 676–777 Ithaca NY: CLC Publications.

The learnability of A-bar chains

Jacqueline van Kampen*

Utrecht University

This paper discusses the learnability of long *wh*-questions. Taking longitudinal data from child Dutch as an example, I will show that the acquisition of long *wh*-movement has been thoroughly prepared in previous acquisition steps. Each step defines a local relation that is preserved in the next acquisition step. The long *wh*-questions appear first with an intermediate pronoun in the speech of the child. My data show that the intermediate pronoun relies on the (previously acquired) relative paradigm which appears in Dutch as an apparently irregular mixture of *d*-pronouns and *w*-pronouns. The present view on the learnability of A-bar chains will lead to the following conclusions: (i) Long *wh*-movement can (hence must) be successive cyclic; (ii) Pied-piping follows from the preservation of licensing contexts only; (iii) Syntactic islands need not to be “learned”. They follow from a non-overlap of local movement domains.

1. Introduction

The study of general grammatical principles is hardly concerned with progress in the study of language acquisition. Yet, child language is the area where limited data already reveal the general principles of grammar. For example, data from the acquisition of A-bar chains (Van Kampen 1997, 2004, 2007) may support the claims about locality in (1).

- (1) Grammar is acquired from the most local patterns possible and the final result preserves much of that locality.
 - a. All movement, e.g. *wh*-movement, is learned in a short step first. Long-distance movements follow from short steps and a fortuitous overlap of initial localities.
 - b. Islands need not be learned. They follow from the fortuitous non-overlap of such localities.

*I would like to thank an anonymous reviewer for valuable and extensive comments. The research for this paper was supported by NWO grant 360-70-290.

The data that will support the claims in (1) are given in (2). The sentences in (2) are long *wh*-questions in child Dutch with an intermediate pronoun, a *w*-pronoun in (2a) and a *d*-pronoun in (2b).

- (2) a. Welke jongen denk je *wie*⟨+animate⟩ daar loopt?
 which boy think you who there walks?
 ‘Which boy do you think is walking over there?’
 b. Welke jongen denk je *die*⟨+gender⟩ ik leuk vind?
 which boy think you that I nice find?
 ‘Which boy do you think I like?’

Long *wh*-questions with an intermediate pronoun appear in the acquisition of various languages (Thornton 1990 for English, Van Kampen 1997 for Dutch, Strik 2006 for French, Gutiérrez 2006 for Basque and Spanish). Although the type with intermediate pronoun is attested in the adult speech for a number of adult languages (see references in Van Kampen 1997: 143), it does not do so in the adult speech of any of the languages just mentioned. Just for that reason, it is remarkable in (2) that child Dutch varies the intermediate pronoun as *die/wie*, whereas adult standard Dutch has no intermediate pronoun and only fits in the neutral complementizer *dat*. The long *wh*-question type in (3) with neutral complementizer was the only maternal input for the children considered here.

- (3) Welke jongen denk je *dat* Sarah leuk vindt?
 which boy think you that Sarah nice find?
 ‘Which boy do you think Sarah likes?’

Hence, we have here the paradoxical fact that child language introduces spontaneously a variation not present in the adult input.

An important circumstance regarding the intermediate pronouns in (2) is that the acquisition of long *wh*-questions in child Dutch follows the earlier acquisition of relative pronouns. Relative pronouns show an agreement pattern with an antecedent. The relative paradigm in Dutch is quite irregular. These irregularities reappear on the intermediate pronoun paradigm. This suggests that the intermediate pronoun is to be analyzed as an ‘A-bar pronoun’ spelled out under local agreement with the moved *wh* phrase (Van Kampen 1997: Chap. 6, cf. Thornton 1990, Thornton & Crain 1994). The variation of the intermediate pronoun in child Dutch (*wat, wie, waar, welk, die, (dat)*) follows the relative agreement paradigm and can be explained from that perspective. Below I will first give a description of the paradigm for A-bar pronouns in Dutch.

1.1 A-bar pronouns

Following Postal (1966), I will label all pronouns as referential indicators ⟨+D⟩. A-bar pronouns are pronouns that must end up in a sentence-initial CP position (Van Kampen

1997: 92ff).¹ They represent pronominal categories with the additional feature $\langle +C \rangle$.² The best examples of inherently A-bar pronouns are *w*-pronouns in root questions and relative pronouns. The V2nd languages have in addition an A-bar pronoun used as a topic in root sentences, the *d*-pronoun, see (4).

- (4) Ze zag daar een jongen. *Die* (= de jongen) vond ze wel leuk
 she saw there a boy. That (= the boy) found she rather nice
 ‘She saw a boy over there. She thought he was rather nice’

The entire paradigm of $\langle +C \rangle$ question pronouns in Dutch is $\langle +wh \rangle$ (*wat*, *wie*, *waar*) and that of $\langle +C \rangle$ topic pronouns is $\langle -wh \rangle$ (*dat*, *die*, *daar*), see the list in (5).

(5) A-bar pronouns in Dutch root sentences

	referent	structural case	oblique case
<i>d</i> -set	$\langle +\text{gender} \rangle$	<i>die</i>	----
	$\langle -\text{gender} \rangle$	<i>dat</i>	([<i>daar</i>] ... (P))
<i>w</i> -set	$\langle +\text{animate} \rangle$	<i>wie</i>	[P <i>wie</i>]
	$\langle -\text{animate} \rangle$	<i>wat</i>	[<i>waar</i>] ... (P)

Agreement is controlled by the $\langle \pm\text{gender} \rangle$ and $\langle \pm\text{animate} \rangle$ properties of the referent.³ The topic *d*-pronouns have an antecedent and they are sensitive to phi-features $\langle \pm\text{gender} \rangle$ of that antecedent. In Dutch, singular nouns can be $\langle +\text{gender} \rangle$ or $\langle -\text{gender} \rangle$. This feature determines the choice of the definite article, either *de vrouw* $\langle +\text{gender} \rangle$ (‘the

1. See for a typology of $\langle +C \rangle$ A-bar pronouns also Rizzi (1990: 67f). I abstract away from Rizzi’s (1997) cartographic approach of the left periphery, which is of no immediate relevance for the discussion.

2. If one allows the category feature $\langle +C \rangle$ to appear in the lexicon as an option for certain pronouns, one may handle for example *wat* as $\langle +D, \pm C \rangle$. The *w*-pronoun *wat* may then appear as indefinite pronoun in $\langle -C \rangle$ argument positions, see (i). As an indefinite argument *wat* cannot rise into the subject position, and remains in situ as in (i). See Postma (1994) (cf. also Cheng 2001).

- (i) a. Als (er) hem *wat* {lukt/bevalt/hindert/tegenzit}
 if (there) him something {succeeds/pleases/bothers/goes against}
 (if he succeeds in something, etc..)
 b. Er is wel *wat* in de keuken
 there is indeed something in the kitchen
 (presumably there is something in the kitchen)

3. I take here $\langle \pm\text{animate} \rangle$ as the relevant feature for $\langle \pm\text{human-like} \rangle$.

woman') or *het huis* ⟨-gender⟩ ('the house').⁴ The plural definite nouns are always *de* (*de vrouwen*, *de huizen*).⁵ Slightly different from the traditional descriptions, I consider number as adding the feature ⟨+gender⟩/*de*. The oblique pronoun *daar* is not sensitive to gender. The question *w*-pronouns have no syntactic antecedent, but they indicate nevertheless whether the answer must be ⟨±animate⟩. The paradigm of the relative pronouns in Dutch is a fixed, but irregular mixture of the forms present in the question *w*-paradigm and the topic *d*-paradigm. The *d*-option for relative pronouns is probably present in V2nd languages only (Van Kampen 2007).

A-bar pronouns are also related to an argument position. The A-bar pronouns in Dutch express structural versus oblique properties. This ⟨±oblique⟩ feature is clearly not related to the antecedent, but to the A(argument)-position, see the examples in (6). Note that *kast* ('cupboard') is a ⟨+gender⟩ *de*-noun.

- (6) a. Op welke kast⟨+oblique⟩ die⟨+gender⟩ jij hebt
 on which cupboard that you have
 t_d gekocht ligt al dat stof?
 t_d bought lies all that dust?
- b. De kast⟨+gender⟩ waar⟨+oblique⟩ al dat stof
 the cupboard where all that dust
 op t_{wh} ligt heb jij gekocht
 on t_{wh} lies have you bought

The general properties of A-bar pronouns are now given in (7).

- (7) *Properties of A-bar pronouns*
- are characterized by ⟨+D⟩, ⟨+C⟩ and ⟨±wh⟩.
 - express phi-features for ⟨±gender⟩ (*d*-pronouns) and ⟨±animate⟩ (*w*-pronouns) of the antecedent/referent.
 - express the ⟨±oblique case⟩ of the argument position.

The intermediate pronouns in long distance questions are A-bar pronouns as well. They have the characteristics in (8).

- (8) a. Their position is a sentence-initial ⟨+C⟩ position.
 b. Their form is partly taken from the root *w*-pronouns ⟨+C, +wh⟩, and partly from the root *d*-pronouns ⟨+C, -wh⟩ in languages that have them.

4. For the binary representation of gender values for nouns in Dutch see Rooryck (2003). Rooryck takes gender to be a univalent feature (cf. Van Kampen 2007), but that is not immediately relevant for the present overview.

5. The indefinite article is always *een* (*vrouw/huis*) irrespective of the gender of the noun. The indefinite plurals have the article \emptyset (*vrouwen/huizen*), again irrespective of gender.

Since the intermediate pronouns show the same variation between *w*-options and *d*-options as the relative pronouns, the *d*-option for intermediate pronouns is probably restricted to V2nd languages as well.⁶

1.2 Organization of the paper

The paper below is organized as follows. Section 2 ('Morphological preliminaries: The relative pronoun paradigm') recapitulates Van Kampen (2007) and exemplifies the irregular variation in the relative paradigm.

Section 3 ('Syntactic preliminaries') settles certain syntactic issues, the acquisition of Ross' (1967) Left Branch Condition in Dutch/German but not in Polish, and the simultaneous acquisition of obligatory pied-piping.

Section 4 ('A-bar pronouns from Spec-head agreement') offers the empirical core of the paper. The extended A-bar chain of a *wh*-phrase first appears in Dutch child language as a chain with the full *wh*-phrase in sentence-initial scope-position and corresponding *w*-set and *d*-set A-bar pronouns in all intermediate positions. See the structure diagram in (9) and the examples in (10).

- (9)
- | | |
|---|--|
| $[\text{Spec}, C_{\langle +wh \rangle} \ C_{\langle +Q \rangle} \ \dots [t_{\langle +wh \rangle} \ C_{\langle +agr \rangle} \ \dots [t_{\langle +wh \rangle} \ C_{\langle +agr \rangle} \ \dots t_{DP\langle +wh \rangle}]_{CP}]_{CP}]_{CP}$ | $\downarrow \qquad \qquad \qquad \downarrow$ |
| <i>welke jongen</i> <i>denk</i> <i>je</i> $t_{\langle +wh \rangle}$ $\left\{ \begin{array}{l} \text{wie} \\ \text{die} \\ \text{dat} \end{array} \right.$ <i>hij zegt</i> $t_{\langle +wh \rangle}$ $\left\{ \begin{array}{l} \text{wie} \\ \text{die} \\ \text{dat} \end{array} \right.$ <i>ik</i> $t_{DP\langle +wh \rangle}$ <i>leuk vind?</i> | |
- (which boy do you think he says I like?)

- (10) a. [Welke jongen] denk je wie daar loopt? (Laura 8;3.8)
 which boy think you who there walks?
 'Which boy do you think is walking over there?'
 b. [Wie] denk je wie er in de auto rijdt? (Laura 8;3.8)
 who think you who there in the car drives?
 'Who do you think drives the car?'
 c. [In welk huis] denk je waar jij woont? (Sarah 4;10.20/
 in which house think you where you live? Laura 7;7.10)
 'In which house do you think that you live?'

6. I assume that the English element *that* in *the man that she looked at* is a (relative) constant (traditionally called 'complementizer') rather than a (relative) pronoun. The same holds for the intermediate *that* in *who do you think that I like?* It may be argued that the 'complementizer' diachronically derives from the most unmarked *d*-pronoun. See Allen (1980) for (relative) *d*-pronouns in old, V2nd, English.

The intermediate A-bar pronouns in C^0 can be derived by the already acquired rule for relative agreement. This leads to an obvious point and a curious prediction. The obvious point is that the intermediate C^0 content is not selected by the matrix verb. It follows from antecedent agreement, not from selection by a matrix verb. The curious prediction runs this way. The irregularities of relative agreement reappear in the intermediate pronouns. This prediction is correct and far from trivial. The Dutch A-bar paradigm for relatives is an intricate mixture of *d*-pronouns and *w*-pronouns (Van Kampen 2007). Significantly, that system has been firmly acquired a year before the chains with intermediate pronouns appear. The evidence that it is the same system that gets active is as striking as the relative paradigm is irregular in Dutch. At the same time, the support for a multiple short step analysis of long *wh*-movement could not be better, since the options that the relative paradigm allows are open at each intermediate point: (*welke jongen C⁰ dat C⁰ die C⁰ wie*), but they are not present in the adult input. In order to acquire the adult system, the child only needs to suppress the spell-out of the agreement and use a complementizer instead, cf. (3). In this view, the top of the chain must remain the most specified element, since it is the antecedent in all intermediate moves. Let me add that, in the corpora considered, this spell-out of the intermediate C^0 , although attested in Dutch dialects (Barbiers, Koenenman & Lekakou 2009), was non-existent in the maternal input, but absolutely dominant and quite long-lasting in the language of the child. This may explain why the type in (9)/(10) pops up at the internet, especially at teenager chat-sites.

Section 5 ('Partial movement') agrees with Fanselow and Mahajan (1996) that German 'partial movement' constructions do not arise from long-distance movement. These constructions are (marginally) attested in Dutch child language as well, see (11).

- (11) *Wat_j denk je t_{(wh)j} wat_i ik t_{(+wh)i} voor liedje zing?* (Sarah 5; 5.12)
 what think you what I for song sing?
 'What kind of song do you think that I sing?'


There are in the Fanselow and Mahajan view two chains in example (11). The *wat* element of the *j*-chain is an expletive pronoun for the subordinate CP. My analysis here follows Fanselow and Mahajan (1996), but I will in addition assume that due to the expletive element, the subordinate CP will be an adjunct of the matrix construction rather than a complement. The matrix expletive is not seen as a reduced form of the complement *wh*-phrase.

Section 6 will offer the learnability perspective. It will be shown that locality is a crucial and inevitable ingredient for the learnability of long-distance relations. The necessity of 'learning from local steps' is supported by the acquisition data. Since the child's acquisition steps show the locality in overt syntax, successive cyclic movement is the best hypothesis for the observed data.

The final Section 7 will state the general conclusions and consider how the study of child language analyzes the learnability of grammatical patterns and why that gives further substance to the broad common sense assertion that grammatical principles must somehow reflect a “psychological reality”.

2. Morphological preliminaries: The relative pronoun paradigm

Relative pronouns are A-bar pronouns that agree in phi-features with the antecedent. The relative pronoun is an A-bar pronoun that relates to the case properties ($\langle \pm \text{oblique} \rangle$) of the argument gap $t_{\langle +C \rangle}$ and to the phi-features of the DP that is the sister of its CP-projection, see (12) (cf. Section 1.1).⁷



The pronominal paradigm of the relative in Dutch is a mixture of *d*-pronouns and *w*-pronouns. If it is possible to express gender agreement with the antecedent, the *d*-pronoun is selected as in (13). Note that the oblique *daar* is not sensitive to gender (cf. Section 1.1) and by consequence *daar* is not selected as a relative.

(13) Dutch relative pronouns with $\langle \pm \text{gender} \rangle$ agreement. The *d*-set comes in: *die*, *dat*, **daar*

	structural case	oblique case
$\langle + \text{gender} \rangle$	<i>die</i>	----
$\langle - \text{gender} \rangle$	<i>dat</i>	*[<i>daar</i>] ... (P)

- a. de jongen *die* $\langle + \text{gender} \rangle$ ik leuk vind
the boy that I like
- b. het huis/het meisje *dat* $\langle - \text{gender} \rangle$ ik leuk vind
the house that I like
- c. *in het huis *daar* $\langle + \text{oblique} \rangle$ ik woon
the house where I live (in)

7. I assume the standard analysis of relative clauses as part of a complex DP.

If it is not possible to express gender agreement, the *w*-pronoun, sensitive to $\langle \pm \text{animate} \rangle$, is selected as in (14). This includes ‘fused’ relatives when there is no antecedent, (14a,b). In contrast to the question *w*-pronoun, cf. (5), the oblique pronoun *waar* can be used with $\langle \pm \text{animate} \rangle$ antecedents, (14c,d). Parallel to the question *w*-pronoun, see (5), only pronouns that are marked for $\langle +\text{animate} \rangle$ (i.e. *wie*) can realize [P pronoun] (14e).

- (14) Dutch relative pronouns with $\langle \pm \text{gender} \rangle$ agreement blocked due to the absence of the antecedent or due to $\langle + \text{oblique} \rangle$. The *w*-set takes over: *wie, wat, waar*

	Structural case	oblique case
⟨+animate⟩	<i>wie</i>	[<i>waar</i>]...P[P <i>wie</i>]
⟨-animate⟩	<i>wat</i>	[<i>waar</i>] ...(P)

- Wie(+animate) ik leuk vind, is het hoofd van de school
Who I like, is the head of the school
- Wat(-animate) ik leuk vind is die bank
what I like, is that couch
- in het huis *waar*(+oblique) ik woon
in the house where I live
- het huis/de jongen *waar*(+oblique) ik verliefd *op* ben
the house/the boy with whom I am in love
- de jongen *op wie*(+animate) ik verliefd ben
the boy with whom I am in love

There are two major exceptions to the generalizations of the selection in (13)–(14). First, (15) shows that the pronoun *die* may be used as ⟨+animate⟩ in relative agreement, although gender agreement (*dat*, cf. (13b)) might have been possible. This parallels with the tendency in spoken Dutch to extend *die* to ⟨–gender⟩ antecedents that are semantically specified for ⟨+animate⟩, see (16). In that case, semantic animacy overrules grammatical gender, which is not perceived by the speaker.

- (15) het meisje<-gender>/<+animate> die<+animate> ik leuk vind
the girl that I nice find
'The girl that I like'

- (16) Neem nou het buurmeisje_{⟨-gender⟩/⟨+animate⟩}. Die_{⟨+animate⟩} vind ik leuk
 take now the neighbors-girl. That find I nice
 'Take the girl of the neighbors. I like her'

Second, (17) shows that the pronoun *wat* may be used as with ⟨-gender⟩ antecedents in relative agreement. This selection of *wat* is preferred in spoken Dutch over the selection of *dat* in (13b).

- (17) het huis/het meisje *wat*⟨-gender⟩ ik leuk vind
 the house/the girl what I like

Both irregularities are analyzed and explained in Van Kampen (2007) as the outcome of a selection problem related to the order of acquisition steps. The acquisition of gender is too slow to suppress the ⟨+animate⟩ agreements of the *w*-system.

The scheme in (18) summarizes the variation in the relative paradigm. The set in (18) covers all observations made in the standard grammar ANS (Algemene Nederlandse Spraakkunst 1997). The irregularity of the relative paradigm is here only stated and exemplified. See for an analysis Van Kampen (2007).


(18)

	<+gender>	<-gender>		no gender	
	gender 'over-rules' animacy	<+animate>	<-animate>	<+animate>	<-animate>
⟨-oblique⟩	de jongen (the boy) <i>die</i> * <i>wie</i>	het meisje (the girl) <i>dat</i> <i>wat</i> <i>die</i> * <i>wie</i>	het paard (the horse) <i>dat</i> <i>wat</i>		
⟨+oblique⟩				de jongen (the boy) <i>waar</i> ... P P <i>wie</i> *P <i>die</i>	het huis (the house) <i>waar</i> ... (P) * <i>daar</i> ... (P)
no antecedent				<i>wie</i>	<i>wat</i> <i>waar</i>

The factual irregularity of the A-bar relative paradigm is what matters here. It will be used as an argument to show that the relative paradigm reappears as a filter on the A-bar agreement in successive cyclic long *wh*-chains (Section 4).

3. Syntactic preliminaries

The analysis of the intermediate A-bar pronouns as resulting from a local Spec-head agreement implies that long *wh*-movement is successive cyclic. It does not explain, though, why the agreement does not appear when the *wh*-antecedent makes its final landing in a Spec.C. Nor has it been explained why, in its first move, the $D_{\langle+wh\rangle}$ obeys the Left Branch Condition and pied-pipes the NP, but why it does not pied-pipe the CP. Why doesn't the Left Branch Condition hold for the ⟨+wh⟩-marked CP? See the structure in (19).



The first move of the *wh*-element *welke* ('which') in (19) pied-pipes (moves along) the NP *jongen* ('boy'). This initial pied-piping continues to be relevant for the subsequent moves of the *wh*-element from Spec.C to Spec.C. By contrast, the subsequent cases of move $\langle +wh \rangle$ do not pied-pipe the CP *wie/die/wat/dat ze kuste* ('that she kissed'). The CP is stranded.

The explanation for this pied-piping asymmetry is quite simple by a re-interpretation of Ross' (1967) Left Branch Condition as proposed in Van Kampen (1997: Chapt. 5, 2004). The original NP complement (*jongen*) of the *wh*-element (*welke*) is pied-piped since as an NP it needs the D° *welke* as a case licenser, see (20).

- (20) a. *welke kuste zij [_D t_{wh} [jongen]_{NP}]_{DP}?
 which kissed she [_D t_{wh} [boy]_{NP}]_{DP}?
 b. [_D welke [jongen]_{NP}]_{DP} kuste zij t_(wh)?
 [_D which [boy]_{NP}]_{DP} kissed she t_(wh)?
 ‘Which boy did she kiss?’

When case targets N' , i.e. in languages with morphological case (Slavic, Latin), the complement NP does not need the D^0 element, the Left Branch Condition does not hold and subextraction of the $D^0_{\langle +wh \rangle}$ is possible. In languages without morphological case on the N^0 (Dutch, German) case targets D^0 (as proposed by Lebeaux 1988: 242f) and pied-piping of the NP complement follows obligatorily. Note that it is not the richness of the morphological case-paradigm, but the target position of the

case-marking (either on D^0 or on N^0) that quite naturally activates or deactivates the Left Branch Condition. In partitive constructions (*combien de livres/wat voor boeken* ‘what for books’) the preposition takes care of the case-licensing and pied-piping becomes an option.

Subsequently, there is an economy conflict between preserving major arguments (by pied-piping) or minimally moving only the $D^0_{\langle +wh \rangle}$ by subextraction. When the Dutch child starts using complex *wh*-phrases, she first moves the $D^0_{\langle +wh \rangle}$ alone, see (21a). Only in a later acquisition step, after the age of 4, the entire *wh*-phrase is pied-piped, see (21b).

- (21) a. Welke wil jij [t_{wh} liedje] zingen? (S. 3;7)
 which want you [t_{wh} song] sing?
 ‘Which song do you want to sing?’
 b. [Welke verhaaltje] wil jij t_{wh} voorlezen? (S. 4;2.8)
 [which story] want you t_{wh} read
 ‘Which story do you want to read?’

Adult Dutch allows subextraction in limited contexts that can be explained along the lines above (Van Kampen 2004).⁸ NP complements, though, are obligatory pied-piped, as in (20b)/(21b). The order of acquisition step now shows that stranding is not learned, it comes for free. What is learned is that the remnant has to be pied-piped given certain licensing conditions, i.e. case-marking for NPs. These licensing conditions have to be learned. The enlarged options for subextraction in child language start in a period when licensing of NP by a D^0 is still optional. This perfectly fits the present proposal.


The same pied-piping story holds for the *wh*-phrase on the left branch, i.e. the left edge, of the CP. If the *wh*-phrase is on the left edge Spec of a CP marked as $\langle +Q \rangle$, it cannot be extracted anymore. It is a $\langle +Q \rangle$ licenser for the $CP_{\langle +Q \rangle}$. Therefore in (22a) the entire CP is topicalized, pied-piping the whole clause *jij kent* (‘you knows’) along with the *wh*-phrase *wie* (‘who’). By contrast, the same *wh*-phrase on the left edge Spec of a $CP_{\langle -Q \rangle}$ will not license that CP, will not activate the Left Branch Condition and must be subextracted, see (22b).

- (22) a. [Wie jij kent] $_{CP\langle +Q \rangle}$ weet ik niet $t_{CP\langle +Q \rangle}$
 [who you knows] $_{CP\langle +Q \rangle}$ know I not $t_{CP\langle +Q \rangle}$
 ‘I don’t know who you are familiar with.’
 b. Wie_i denk je [$t_{\langle wh \rangle i}$ [dat ik ken $t_{\langle wh \rangle i}$]] $_{CP\langle -Q \rangle}$?
 who_i think you [$t_{\langle wh \rangle i}$ [that I know $t_{\langle wh \rangle i}$]] $_{CP\langle -Q \rangle}$?
 ‘Who do you think I am familiar with?’

8. The stranding possibilities in Dutch hold for prepositions with an *r*-marked *w*-/*d*-pronoun (*waar/daar*) and for AP complements of a $Deg^0_{\langle +wh \rangle}$. See Corver (1990: 195ff).

This shifts the problem. It may be that the *wh*-phrase in (22b) can be freely extracted, but why did it ever land in such an intermediate Spec.C? After all, the *wh*-movement into an intermediate Spec.C position cannot have been triggered by a target position $C_{\langle +Q \rangle}$, since the weak assertive *denken* selects a $\langle -Q \rangle$ complement. See for this “triggering problem” Lasnik and Saito (1984). There is a semantically relevant trigger $\langle +Q \rangle$, but where is the local trigger given the $C_{\langle -Q \rangle}$ in (22b)? My proposal runs as follows. There are two triggers, $\langle +C \rangle$ and $\langle +Q \rangle$. The trigger $\langle +C \rangle$ requires that each A-bar pronoun $\langle +D, +C \rangle$ gets positioned in the first A-bar position beyond the predicate-argument structure. If the D^0 is at the same time a (case) licenser, it will pied-pipe its complement. This may seem a re-description of the contention that rules have to be local, but the position of the first C^0 -up is a clausal scope-position that has to be checked anyway as a $C^0_{\langle +Q \rangle}$ or a $C^0_{\langle -Q \rangle}$ in order to find out whether the *wh*-movement has to be local or (pro)long(ed). If the first C^0 -up is a $C_{\langle +Q \rangle}$, the movement triggering feature $\langle +C \rangle$ will be deactivated, say removed. Otherwise, when the first C^0 -up is a neutral head $C^0_{\langle -Q \rangle}$, like the C^0 of the weak assertive *denken* (‘think’) in (19)/(22b), the movement triggering feature $\langle +C \rangle$ of the *wh*-phrase will not be removed and remain active.

Suppose there is this local movement to the first C^0 -up position, see (23).



The CP complement is an argument of the matrix CP and it is $C^0_{\langle -Q \rangle}$. The left edge of that CP argument is accessible and (after the first *wh*-movement) spotted as a $\langle +C \rangle$. Hence, the $\langle +C \rangle$ /A-bar operation reapplies to the CP argument. There is a target/source overlap in Spec.C. Yet, this time the *wh*-pronoun will not pied-pipe its complement (the CP complement), since it does not license that $CP_{\langle -Q \rangle}$ complement. My central point in this usual explanation will be that all contributing factors in the reapplication of *wh* movement have already been acquired by the child from more elementary constructions. These are (i) accessibility of information at the left edge, (ii) the movement up to the first A-bar position, (iii) the \pm pied-piping factor, (iv) the A-bar agreement from the relative paradigm (as will be shown in Section 4).

My point is of course not to re-describe Chomsky's (1973) Spec.C escape hatch or Ross' (1967) Left Branch Condition and pied-piping. I only argue that the ingredients (i) to (iv) for long *wh*-movement are already in place long before the child moves on to apply them in a combined fashion. So, my point is that long *wh*-movement is not learned, but follows as an implication from simple acquisition steps. Not a priori and innate, but previously discovered in more elementary and highly frequent structures.

4. A-bar pronouns from Spec-head agreement

Chomsky (1973) argues that the long *wh*-movement had to be the outcome of a successive cyclic passage of the *wh*-phrase through the intermediate Spec.C positions.⁹ This successive cyclic passage is reflected in child Dutch as an appearance of A-bar pronouns in the intermediate C⁰ head positions. It seems natural to derive these intermediate A-bar pronouns by means of a Spec.C-C⁰ agreement as in Thornton (1990), Van Kampen (1997: Chapt. 6). The chain is formed by the Spec.Cs, whereas the spell-out of the C⁰s is a reflection of local agreement. For that reason, the C⁰s do not form a chain, cf. the structure in (9)/(19).

The following point is of special interest. Standard adult language evades the use of A-bar pronouns in the intermediate positions of long *wh*-movement. It restricts itself to the neutral C⁰ head (complementizer) *dat* only. Child language, by contrast, applies the intermediate A-bar pronouns for a long time almost exclusively. See the numbers in (24) for Laura and Sarah.

(24)	agreeing pronoun	complementizer <i>dat</i>
Laura (between 7;2–9;0)	74	5
Sarah (between 4;7–6;0)	12	0

Two things must be stressed here. First, long-distance questions appear quite late. Sarah's first long-distance questions are attested after the age of five. There are no long-distance questions attested in the speech of Laura before the age of seven. This is long after the paradigm of question *w*-pronouns and especially the paradigm of relative

9. See Boeckx (2003) about the short-move/long-move issue and why a series of short moves may count as more "economical" than a single "suitably delayed" long movement. Boeckx (2003) argues, a bit circularly, that the delay for a single long movement requires the same amount of repetitive structure checking. The checking procedure is so to speak a bit shortsighted. It first spots C⁰, but then it needs a closer local inspection to see the C⁰_(+Q) or the C⁰_(-Q) specification. If one assumes that the multiple short steps can do with a local checking without restarting the cyclic machine, it is a decisive advantage. The present acquisitional argument is meant to be empirical rather than conceptual.

agreement have been firmly established in the speech of the child. Second, although the intermediate A-bar pronouns appear spontaneously, this does not mean that they have to be learned or that their pattern is innate.

Below I will give all possible and attested variants of the agreeing pronoun. The rule of relative agreement seems to apply in all relevant cases, as I will show.

Simple (non-complex) *w*-pronouns don't have an N-complement. The agreeing A-bar pronoun in intermediate position needs only to vary for $\langle \pm \text{animate} \rangle$ and $\langle \pm \text{oblique} \rangle$, not for $\langle \pm \text{gender} \rangle$ properties of the antecedent. In (25) all examples that should be possible are given and (26) lists the corresponding attested examples in child Dutch (Van Kampen corpus CHILDES and diary notes 1993–1997).

- (25) a. Wie $\langle +\text{animate} \rangle$ denk je *wie* $\langle +\text{animate} \rangle$ ik leuk vind?
'Who do you think I like?'
b. Wat $\langle -\text{animate} \rangle$ denk je *wat* $\langle -\text{animate} \rangle$ ik leuk vind?
'What do you think I like?'
c. Waar $\langle +\text{oblique} \rangle$ denk je *waar* $\langle +\text{oblique} \rangle$ ik woon?
'Where do you think I live?'
(26) a. Wie denk je *wie* er in de auto rijdt? (Laura 8;3.8)
who think you who there in the car drives?
'Who do you think drives the car?'
d. Wat denk je *wat* ik ga zeggen? (Sarah 6;4.13)
What think you what I go say?
'What do you think I will say?'
c. Waar denk je *waar* mijn handen zijn? (Sarah 4;10.20)
where think you where my hands are?
'Where do you think my hands are?'

The set of agreeing *w*-pronouns in (25)–(26) is not complete. Example (27) occurs as well.¹⁰

- (27) Wie denk je *die* er jarig is? (Laura 9;1.4)
Who think you that there 'jarig' is?
'Who do you think has her/his birthday?'

In Section 2 it was argued that spoken Dutch has the tendency to select *die* with antecedents that are semantically specified for $\langle +\text{animate} \rangle$, cf. (16). It was shown to hold for


10. Van Kampen (1997: appendices B and C) reports the spell-out of intermediate attributive *welke* and *w*-adverbs like *hoe*, *wanneer* and *waarom*. I will restrict the present analysis to non-attributive pronouns.

relative agreement as well, cf. (15). I assume *die* in (27) to be specified for $\langle +animate \rangle$ as well.


The picture is a bit more complicated when full, complex, *wh*-phrases exhibit long successive movement. In fact, the agreement properties known from the relative paradigm, with all alternatives and irregularities, seem to reappear in the paradigm of the intermediate A-bar pronoun. These properties were already learned in a previous acquisition step.

The CP-adjoined DP closes further grammatical calculations in the CP with an agreement that checks the carry-over of information. Like the relative pronoun in (28a) (cf. (12)), the intermediate agreeing pronoun in (28b) agrees in phi-features with the moved *wh*-phrase.

(28) a. relative pronoun



b. intermediate pronoun



There is a difference between the top-labels in (28a) and in (28b) and the phrase “CP-adjoined DP” covers up that difference. It appears to be irrelevant for the A-bar agreement.

In (29) all possibilities of intermediate pronouns agreeing with a complex *wh* phrase are given. The numbers between brackets refer to the parallel relative pronouns in Section 2. The variation of the intermediate pronoun shows that their choice is due to agreement, not to movement.¹¹

11. This contrasts with Barbiers, Koenenman and Lekakou (2009). Also in contrast to Barbiers, Koenenman and Lekakou (2009) is my assumption that definiteness does not play a role for the $\langle \pm wh \rangle$ A-bar pronoun. I do not see in which case $\langle \pm definite \rangle$ really proves to be a phi-feature. Compare the following contrastive minimal pairs (cf. (26a)/(27)), where *er* (‘there’) signals

- (29) a. Welke jongen denk je *die*(+gender) ik leuk vind?
 ‘Which boy do you think I like?’ (cf. (13a))
- b. welk huis/welk meisje denk je *dat*(-gender)dat ik leuk vind?
 ‘Which house/which girl do you think I like?’ (cf. (13b))
- c. in welk huis denk je *waar*(+oblique) ik woon?
 ‘In which house do you think I live?’ (cf. (14c))
- d. op welk huis/welke jongen denk je *waar*(+oblique) ik *op* verliefd ben?
 ‘Which house/which boy do you think with whom I am in love’ (cf. (14d))
- e. Op welke jongen denk je *op wie*(+animate) ik verliefd ben?
 ‘Which boy do you think with whom I am in love’ (cf. (14e))
- f. Welk meisje denk je *die*(+animate) ik leuk vind?
 ‘Which girl do you think I like?’ (cf. (15))
- g. welk huis/welk meisje denk je *wat*(-gender) ik leuk vind?
 ‘Which house/which girl do you think I like?’ (cf. (17))

The parallel attested examples in child Dutch are given in (30). Example (30b) with *dat* will further be left out of the discussion. It is a neutral C⁰ head generalized to all long distance questions in standard Dutch. Moreover it is not or hardly attested in the child data, cf. the figures in (24).

- (30) a. Welke stift(+gender) denk je *die* ik moet nemen? (Laura 7;8.18)
 which felt-tip think you that I must take?
 ‘Which felt-tip do you think I must take?’
- b. Welk meisje denk je *dat* ik leuk vind (adult Dutch)
 ‘Which girl do you think I like?’
- c. In welk huis denk je *waar* jij woont? (Sarah 4;10.20/
 in which house think you where you live? (Laura 7;7.10)
 ‘In which house do you think you live?’

the ⟨-definite⟩ subject. Removing *er* from the relative examples (ii) makes the sentences grammatical.

(i) Intermediates

- a. Wie denk je *die* t_{wh} er jarig is?
 b. Wie denk je *wie* t_{wh} er jarig is?
 who think you that t_{wh} there ‘jarig’ is?
 ‘Who do you think has her/his birthday?’

(ii) Relatives

- a. *de jongen [*die* er jarig is] (the boy that has his birthday)
 b. *het meisje [*wat* er jarig is] (the girl that has her birthday)

- d. Op welke stoel denk je *waar* ik op zit? (Laura 7;10.1)
 on which chair think you where I on sit?
 'On which chair do you think I sit?'
 e. Not attested, but possible. Attested is:
 Op wie denk je *op wie* Sarah verliefd is? (Laura 8;3.8)
 with who think you with who Sarah in love is?
 'With whom do you think Sarah is in love?'
 f. Welk meisje denk je die ik een hand geef? (Laura 8;3.8)
 which girl think you that I a hand give?
 'Which girl do you think I shake hands with?'
 g. Welk cadeau denk je *wat* ik geef? (Laura 7;10.1)
 which present think you wat I give?
 'Which present do you think I will give?'

The relative agreement paradigm successfully excludes the cases in (31). These are unattested in child language.

- (31) a. *in welk huis denk je *daar*<+oblique> ik woon?
 in which house think you I live? (cf. (13c))
 b. *wat voor boeken<+plural> denk je *wat* ik heb gelezen?
 what for books do you think what I have read?
 what kind of books do you think I have read?
 c. *welke boeken<+plural> denk je *wat* ik heb gelezen?
 which books do you think what I have read?
 which books do you think I have read?
 d. *welke villa<+gender> denk je *wat* ik ga kopen
 which villa do you think I will buy?

All intermediate obliques have to be <+wh> *waar* and cannot be <-wh> *daar*, as already predicted by relative agreement, cf. (13c). The A-bar pronoun *wat* cannot agree with the plural *boeken* in relative agreement nor can it be spell-out of agreement in (31b,c). The same holds for *wat* in (31d). It cannot agree with the <+gender> noun *villa*. As a relative pronoun *wat* can only refer to a <-gender> noun. Therefore, *wat* can agree with the <-gender> noun *spelletje* in (32), cf. (30g) and (17).

- (32) a. Wat voor spelletje<-gender> denk je *wat* ik
 what for game think you what I
 wil doen ? (Laura 7;9.27)
 want do ?
 'What kind of game do you think I want to do?'

The agreement solution also predicts (correctly) that the intermediate positions are never filled in by a repetition of the *wh*-phrase. The intermediates are for pronominal forms only. See the (unattested) examples in (33).

- (33) a. *welke jongen denk je welke jongen ik leuk vind?
 which boy think you which boy I like?
 b. *welk huis denk je welk huis ik leuk vind?
 which house think you which house I like?
 c. *in welk huis denk je in welk huis ik woon?
 in which house think you in which house I live?

The agreement rule may extend to the preposition of oblique phrases, cf. (30e). If the pronoun can express inherent case, as in (34a,b) (cf. (30c)), the preposition is not repeated. In (34a) *waar* reflects the locative. It corresponds with the antecedent *op welke school* ('at which school'). In (34b) *wie* is inherently marked for dative ('mee-werkend voorwerp'), which is possible in Dutch (ANS 1997: par. 5.5.8.2). However, if the intermediate *w*-pronoun cannot reflect oblique case, the preposition is repeated. In (34c) *wie* would not be a replacement of *op wie*, cf. (14e) (Van Kampen 1997: 151f).

- (34) a. Op welke school denk je waar Laura zit? (Laura 8;3.8)
 at which school think you where Laura sits?
 'At which school do you think Laura is?'
 b. Aan wie denk je wie ik een brief schrijf? (Laura 7;9.2)
 to who think you who I a letter write?
 'To whom do you think I write a letter?'
 c. Op wie denk je [op wie]_{Co} Sarah verliefd is? (Laura 7;10.25)
 with who think you with who Sarah in love is?
 'With whom do you think Sarah is in love?'

The present approach suggests that the oblique preposition and its A-bar pronoun fit into the C⁰ head position as a complex head. See the brackets in (34c) above for the anomalous analysis. [P + A-bar features]_{Co}. The nice outcome of the present approach is that it successfully explains by relative agreement the grammaticality of (34c) versus the ungrammaticality of the three examples in (33).¹²

This leaves me with a final difficulty. The intermediate A-bar pronoun *wie* is correct as a ⟨+animate⟩ spell-out of an intermediate A-bar pronoun for the phrase *welke jongen* and *welk meisje* in (35). Yet, it is not predicted by relative agreement which would spell out the also correct *die* for ⟨+gender⟩/⟨+animate⟩ antecedents, see the scheme in (18) (cf. (13a)/(15)).

- (35) a. Welke jongen denk je wie daar loopt? (Laura 8;3.8)
 which boy think you who there walks?
 'Which boy do you think is walking over there?'
 b. Welk meisje denk je wie ik een hand geef? (Laura 8;3.8)
 which girl think you who I a hand give?
 'Which girl do you think I shake hands with?'

12. See for a different analysis Pankau (2007).

The main rule for relative pronouns is to select the *d*-pronoun if gender agreement is possible and to select the *w*-pronoun otherwise.¹³ Around the age of five when the long *wh*-movements and their intermediate A-bar pronouns begin to appear with some regularity in the speech of the child, the relative agreement pattern is already well established. The relative paradigm reappears for the intermediate pronoun agreement. Yet, the more specific gender agreement that determines the *d*-set weakens to an option under the more complex calculation of long *wh*-movement. Descriptively, the relative A-bar paradigm and the intermediate A-bar paradigm can now be stated as in (36).

- (36) A-bar agreement with a locally adjacent antecedent holds for relative pronouns and for the intermediate pronouns in long *wh*-movements.
- a. Relatives
Select a *d*-set pronoun if the antecedent has gender. Select a *w*-set pronoun otherwise.
 - b. Intermediates
As relatives, or select a *w*-set pronoun if the antecedent is ⟨+animate⟩.

The prediction in (36) seems an excellent underlining that the order of acquisition steps is crucial to understand the learnability of grammar.

From a somewhat broader point of view, one may notice that the present agreement proposal fits with Rizzi's (1996) *wh*-criterion. It also tallies well with the "doubly-filled Comp filter" (Chomsky and Lasnik 1977). Either the C^0 gets realized or the Spec.C, but not both. The trigger ⟨+C⟩ and ⟨+Q⟩ features on the C^0 head are erased when the specifier moves in the final landing side. The non-trigger ⟨−Q⟩ feature attracts phi-agreement. The *wh*-movement feature ⟨+C⟩ is not erased and consequently the *wh*-phrase ⟨+D, +C⟩ moves further.¹⁴

5. Partial movement: Single long (successive) chain or double short chain?

Child Dutch exhibits another type of complex question, see (39).¹⁵

13. Diachronically, Afrikaans and English show that indeed the *w*-set takes over when gender disappears.

14. The doubly-filled Comp filter has also been reinterpreted by Koopman (2000) in a different way and perspective.

15. The partial movement type is also attested in Dutch dialects (Barbiers, Koenenman & Lekakou 2009) and in informal Dutch, see (i). It was not attested, though, in the maternal input of the children considered here.

- (39) a. *Wat denk je waar/wie/wat dit is?*
 what think you where/who/what this is?
 'Where/who/what do you think this is?'

This type of complex question is also attested in e.g. adult German (Fanselow 2006) and labeled 'partial movement'. The *wh*-elements in the partial movement construction do not fit the conception of long-distance string that has been studied above. The *wh* element *wat* in the matrix clause of (39) cannot be changed into *wie* or *waar*, see (40).

- (40) a. *in welk huis/waar denk je waar ik woon?* (long *wh*-movement)
wat denk je in welk huis/waar ik woon? (partial *wh*-movement)
 **waar denk je in welk huis ik woon?* (*partial *wh*-movement)
 in which house do you think I live
- b. *welke jongen/wie denk je wie ik heb gekust?* (long *wh*-movement)
wat denk je welke jongen/wie ik heb gekust? (partial *wh*-movement)
 **wie denk je welke jongen ik heb gekust?* (*partial *wh*-movement)
 which boy/who do you think I kissed?

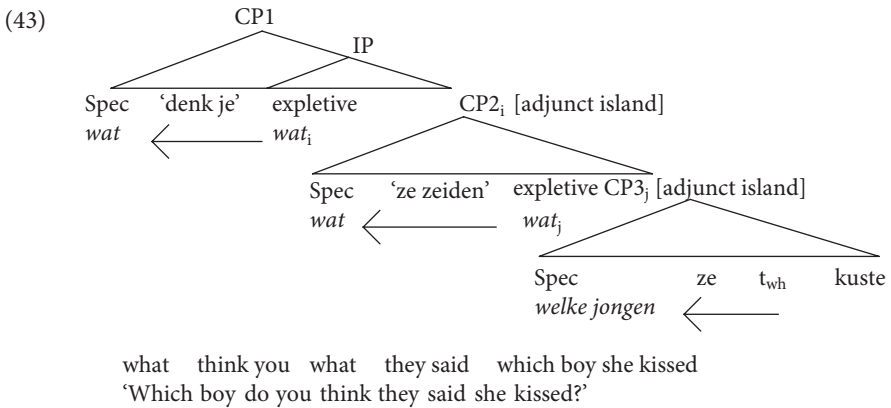
No such invariability for the sentence initial *wh*-phrase was present in any of the long successive *wh*-chains in Section 4. Another difference is that the second *wh*-element in (39), the one in front of the subordinate clause can be expanded in a full *wh*-phrase, see (41) in comparison with (33). Again, (42) states some attested examples in child Dutch.

- (41) a. *Wat denk je [welke jongen] ik leuk vind?*
 what think you which boy I nice find?
 'Which boy do you think I like?'
- b. *Wat denk je [in welk huis] ik woon?*
 what think you in which house I live?
- (42) a. *Wat denk je wat voor woord dit is?* (Laura 10;10.24)
wat think you what for word this is?
 'What kind of word do you think that this is?'
- b. *Wat denk je wat ik voor liedje zing?* (Sarah 5;5.12)
wat think you what I for song sing?
 'What kind of song do you think that I sing?'
- c. *Wat denk je bij de hoeveelste ik ben?* (Laura 7;11.8)
wat think you at the how maniest I am?
 'At which number do you think I am?'

-
- (i) (Irene Moors, TV presenter. On screen 17-12-2007)

Wat denk je waar we naar toe gaan?
 what think you where we at to go?
 'Where do you think that we will go?'

The full *wh*-phrases in the middle of (41) and (42) must be on top of a *wh*-chain of their own. The original notion *wh*-chain, and its explanatory potential, is saved if we assume that the *wat*-constructions in (41)–(42) are based on two chains. For that reason, I accept the proposal by Fanselow and Mahajan (1996, partly following Dayal 1994) that the first *wh*-element *wat* in (41)–(42) is a type of sentential expletive linked to the subordinate $CP_{\langle +Q \rangle / \langle +wh \rangle}$. No long distance chain is involved in partial *wh*-movement. Compare the structure in (21) for long successive *wh*-movement to one in (43) with partial short movements.



This orientation differs from and is incompatible with the *wh*-chain analysis offered in Barbiers, Koenenman and Lekakou (2009). The abstract expletive in (43) functions as a pronominal stand in for the subordinate $CP_{\langle +Q \rangle}$ as a whole. It moves from the matrix object position to the matrix Spec.C, because of the $\langle +Q \rangle$ feature. Its expletive nature explains its inability to change form and its inability to turn into a *wh*-phrase. The subordinate $CP_{\langle +Q \rangle}$ can no longer be in argument position. It is forced to become an adjunct to the expletive, because its complement function and object theta relations are taken over by the expletive.¹⁶ Compare the adjunct island CP (with expletive *er*) in (44a) to the complement CP in (44b).

16. Fanselow & Mahajan (1996) argue, on different grounds, that the $CP_{\langle +Q \rangle}$ is a complement. I won't go into the different analyses of partial movement here. See Fanselow (2006) for an overview.

- (44) a. *Welke jongen rekende jij er op
 which boy counted you there on
 [dat je nog t_{wh} kende]_{CPadjunct} ?
 [that you still t_{wh} knew]_{CPadjunct} ?
- b. Welke jongen dacht jij [dat je nog t_{wh} kende]_{CPcomplement} ?
 which boy thought you [that you still t_{wh} knew]_{CPcomplement} ?
 ‘Which boy did you think that you still knew?’

This explains why the verbs *denken* (‘think’) and *zeggen* (‘say’), that do in general not support sentential complements $CP_{\langle +Q \rangle}$ (cf. Section 3), are in fact construed with such a $\langle +wh \rangle$ $CP_{\langle +Q \rangle}$ in (43) (*wat ze zeiden, welke jongen ze kuste*). It is the abstract expletive *wat* that satisfies the insertion frame of the verbs *denken* and *zeggen*. Due to that same expletive, the $CP_{\langle +Q \rangle}$ is in adjunct position and hence compatible with *denken* and *zeggen*.

6. Conclusion: Long-distance movement as an overlap of (cyclic) localities

In this paper I have defended the idea that all movement can be learned from example sets with the shortest steps possible, assuming that long-distance movement and island constraints follow from (i) left edge accessibility for the $\langle +D \rangle$ features, (ii) the minimality condition on movements $\langle +C, +D \rangle$ and (iii) (lack of) pied-piping.

In Section 3 it was argued that the child begins with subextracting the $D^0_{\langle +wh \rangle}$ of a complex *wh*-phrase, as in (45a). In a second step only, the NP complement is pied-piped, as in (45b).

- (45) a. Welke wil jij [t_{wh} liedje] zingen? (S. 3;7)
 which want you [t_{wh} song] sing?
 ‘Which song do you want to sing?’
- b. [Welke verhaaltje] wil jij t_{wh} voorlezen? (S. 4;2.8)
 [which story] want you t_{wh} read
 ‘Which story do you want to read?’


This order of acquisition step shows that stranding is not learned, it comes for free. Pied-piping is learned. Dutch allows subextraction in limited contexts. Hence, the child has to acquire the licensing conditions for the remnant, such as an obligatory D^0 context for case-assignment to NPs.

Subsequently, the child starts using long-distance *wh*-questions. In contrast to simple A-bar pronoun constructions, long *wh*-questions are fairly rare in the input.¹⁷

17. In the Van Kampen corpus (120 files of 45 minutes recordings, a total of 61.526 input sentences) only 4 long-distance *wh*-questions appeared in the speech of the mother.

One cannot say that the child is “trained” on such structures. Nevertheless, they appear after the age of 5 and they take forms not present in the input. Four devices are to be combined by the child, see (46).

- (46)
- The short step trigger $\langle +C \rangle / \langle +Q \rangle$ to $C_{\langle -Q \rangle}$.
 - The left edge accessibility (Chomsky’s 1973 “escape hatch”).
 - Obligatory pied-piping the NP complement.
 - The intermediate A-bar pronoun agreement.



A closer look reveals that all four devices in (46) are known from previous acquisition steps and have been acquired earlier from highly frequent simple contexts, see (47). The ages are a rough indication. Note that long-distance questions appear quite late in the speech of the child (after the age of 5). Long distance questions solicit the hearer’s opinion and are by consequence dependent upon a theory of mind reading (Van Kampen 1997: 141). The young child is a formidable mind-reader, but making the systematic distinction between the inner and outer domain is a different matter and the corresponding grammatical devices come in later.

(47) Already learned

- A-bar pronouns $\langle +C \rangle$ appear in initial position and they leave an argument gap: learned from all simple topic and question sentences. (before the age of 3)
- Left branch subextraction: learned from all *wh*-phrases. (before the age of 4)
- Obligatory pied-piping the NP. (before the age of 5)
- A-bar pronoun agreement: learned from relatives. (between 4–5)

Movement affects the left edge for reasons that were already known from the stranding constructions in (45a) above. So, movement of the *wh*-element and stranding the remaining constituent is old. The fact that the CP does not pied-pipe is old, since, in contrast to NP, the $CP_{\langle -Q \rangle}$ does not need a $\langle +Q \rangle$ licenser and it is only obligatory licensing that forces the learner into pied-piping.

The long *wh*-movements show how a short *wh*-movement to the nearest $CP_{\langle -Q \rangle}$ brought the A-bar pronoun into a new domain for a new short step. There happened to be an overlap of movement localities. This implies that lack of a locality overlap causes an island constraint for *wh*-movement. That seems trivial, but the logical consequence is that syntactic islands need not be learned. They follow from any stupid non-overlap of movement localities. The learner discovers transparency as something already present in the system, the left branch extraction and the pied-piping triggers. The long movements follow without need to notice or learn island constraints.

7. Epilogue

In generative grammar one generally assumes that grammatical distinctions reflect a psychological reality. However, such a statement says very little. It is more a common sense perspective about future developments. This is not to defend skepticism about the perspective as such. It is only to point out that the perspective is as obvious as it is still open and largely unexplored.

On the sunny side of things, language acquisition seems to me to be the first field where mere grammatical distinctions may predict psychologically real performance in some systematic fashion. A basic point is that grammatical categories and principles are acquired in highly local and highly repetitive patterns. That locality, and hence the learnability of the system, is preserved in the final state. Locality in grammar does not seem to be some mysterious innate property, but rather a property imposed by the need to maintain learnability for toddlers.

The more specific point from the acquisition analysis above has been that the previous acquisition of the paradigm for relative A-bar pronouns was crucial to grasp long *wh*-movements. The explicit marking of that *wh*-chain by agreement is spontaneous child language. One cannot say that the long *wh*-movement is learned. In the first place, the long *wh*-constructions are rare in the input. In the second place, they appear in child language almost exclusively with a spontaneous *wh*-agreement pattern. Nevertheless, it would be wrong to conclude from the spontaneous appearance of an A-bar agreement that principles and categories are innate. It rather shows how acquisition works as a recombination of devices already acquired from simplified highly repetitive patterns. The development of a *wh*-chain unmarked by agreement follows later.

An advantage of language acquisition over the analysis of language perception and language production is that the acquisition procedure is one of slow motion. It tracks developments that can be measured in weeks and months rather than in milliseconds. The acquisition performance is like historical change. It can be reconstructed by a typical linguistic method, the study of language corpora. When considering longitudinal acquisition data, ambitious hypotheses, like for example the proposal that Dutch/German is basically SVO (Kayne 1994) runs into considerable problems, but successive cyclicity of *wh*-movement fits the data wonderfully.

References

- Allen, C. 1980. *Topics in Diachronic English Syntax*. New York NY: Garland.
- Barbiers, S., Koenenman, O. & Lekakou, M. 2009. Syntactic doubling and the structure of *wh*-chains. *Journal of Linguistics*.
- Boeckx, C. 2003. *Islands and Chains* [Linguistik Aktuell/Linguistics Today 63]. Amsterdam: John Benjamins.
- Cheng, L. 2001. Het is allemaal *wat*. Inaugural address, Leiden University.

- Chomsky, N. 1973. Conditions on transformations. In *A Festschrift for Morris Halle*, S. Anderson & P. Kiparsky (eds), 232–286. New York NY: Holt, Rinehart and Winston.
- Chomsky, N. & Lasnik, H. 1977. Filters and control. *Linguistic Inquiry* 8(3): 425–504.
- Corver, N.F.M. 1990. The Syntax of Left-Branch Extractions. Ph.D. dissertation, University of Tilburg.
- Dayal, V. 1994. Scope marking as indirect *wh*-dependency. *Natural Language Semantics* 2: 137–170.
- Fanselow, G. 2006. Partial movement. In *The Blackwell Companion to Syntax*, M. Everaert, R. Goedemans, B. Hollebrandse & H. van Riemsdijk (eds). Oxford: Blackwell.
- Fanselow, G. & Mahajan, A. 1996. Partial movement and successive cyclicity. In *Papers on Wh-scope Marking* [Working Papers of the SFB 340], U. Lutz & G. Müller (eds), 131–161. Tübingen: University of Tübingen.
- Geerts, G. et al. (eds). 1997. *ANS* (Algemene Nederlandse Spraakkunst), second revised edn. Groningen: Wolters-Noordhoff.
- Gutiérrez, M.J. 2006. Acquiring long-distance *wh*-questions in L1 Spanish. In *The Acquisition of Syntax in Romance Languages* [Language Acquisition & Language Disorders 41], V. Torrens & L. Escobar (eds), 251–289. Amsterdam: John Benjamins.
- Kampen, J. van. 1997. *First Steps in Wh-movement*. Delft: Eburon.
- Kampen, J. van. 2004. An acquisitional view on optionality. *Lingua* 114: 1133–1146.
- Kampen, J. van. 2007. Relative agreement in Dutch. In *Linguistics in the Netherlands 2007*, B. Los & M. van Koppen (eds), 112–125. Amsterdam: John Benjamins.
- Kayne, R. 1994. *The Antisymmetry of Syntax*. Cambridge Ma: The MIT Press.
- Koopman, H. 2000. *The Syntax of Specifiers and Heads* [Leading Linguist Series 3]. London: Routledge.
- Lasnik, H. & Saito, M. 1984. On the nature of proper government. *Linguistic Inquiry* 15(2): 235–289.
- Lebeaux, D. 1988. Language Acquisition and the Form of Grammar. Ph.D. dissertation, University of Massachusetts.
- Pankau, A. 2007. Die *wh*-Kopie-Konstruktion im Deutschen. MA thesis, University of Potsdam.
- Postal, P. 1966. On so-called pronouns in English. In *Monograph Series on Language and Linguistics* 19, F. Dinneen (ed.), 177–206. Washington DC: Georgetown University Press.
- Postma, G. 1994. The indefinite reading of WH. In *Linguistics in the Netherlands 1994*, R. Bok-Bennema & C. Cremers (eds), 187–199. Amsterdam: John Benjamins.
- Rizzi, L. 1990. *Relativized Minimality*. Cambridge MA: MIT Press.
- Rizzi, L. 1996. Residual verb second and the *wh* criterion. In *Parameters and Functional Heads*, A. Belletti & L. Rizzi (eds), 63–90. Oxford: OUP.
- Rizzi, L. 1997. The fine structure of the left periphery. In *Elements of Grammar*, L.M. Haegeman (ed), 281–337. Dordrecht: Kluwer.
- Rooryck, J. 2003. The morphosyntactic structure of articles and pronouns in Dutch. In *Germania et alia. A Linguistic Webschrift for Hans den Besten*, J. Koster & H. van Riemsdijk (eds). <http://odur.let.rug.nl/~koster/DenBesten/contents.htm>.
- Ross, J.R. 1967. Constraints on Variables in Syntax. Ph.D. dissertation, MIT.
- Strik, N. 2006. L'acquisition des phrases interrogatives chez les enfants francophones. *Psychologie Française* 52: 27–39.
- Thornton, R. 1990. Adventures in Long Distance Moving: The Acquisition of Complex *Wh*-movement. Ph.D. dissertation, University of Connecticut, Storrs.
- Thornton, R. & S. Crain. 1994. Successive cyclic movement. In *Language Acquisition Studies in Generative Grammar*, Hoekstra and Schwartz (eds.), 215–252. Amsterdam: John Benjamins.

Looking at anaphora

The psychological reality of the Primitives of Binding model

Arnout W. Koornneef
Utrecht University

The Primitives of Binding model (Reuland 2001) presents a linguistic framework to explain the complex pattern of anaphora (reflexives and pronominals) across a large number of languages, i.e. it has theoretical validity. The model consists of a syntactic, semantic and discourse module and, furthermore, incorporates an economy hierarchy: syntactic anaphoric dependencies require less processing resources than semantic anaphoric dependencies, which in turn are cheaper than discourse anaphoric dependencies. I will discuss the results of three eye-tracking experiments which suggest that this economy hierarchy is reflected in real time anaphora comprehension. Hence, the model has psychological reality as well. The implications of this conclusion will be illustrated by comparing a real time version of the model to Friederici's (2002) neurocognitive approach to language comprehension.

1. Introduction

One of the primary objectives of theoretical linguistics is to discover what rules and principles constitute the fundamentals of human language. Although they have come a long way, linguists traditionally have not been particularly interested in how our brain represents or uses these rules while processing language expressions in real time. This led to a split between theoretical linguistics and psycholinguistics in the 1970s. The bad news is that from that day on, the gap between them only got bigger, despite the fact that both have the same goal, namely understanding human language. To an objective observer this may come as a big surprise, to say the least, as it seems obvious that the disciplines could benefit substantially from working closely together. Fortunately, the good news is that recent developments point exactly in that direction, that is, theoretical linguists increasingly adopt psycholinguistic methods to put their formal theories to the behavioral and neuro-imaging test and psycholinguists have become aware of the fact that one cannot discard linguistic theories while exploring language processing in real time (see for example Sturt 2003; Burkhardt 2005; Grodzinsky & Friederici 2006;

Runner et al. 2006; Vasic 2006). In other words, researchers from the different fields seem to acknowledge that theoretical validity and psychological reality are inextricably intertwined: one cannot take the position of having a theoretically valid approach while none of its predictions are reflected in real time language comprehension and, vice versa, the value of behavioral and neuro-imaging results diminishes significantly when interpreted in isolation without any reference to well-established linguistic facts. In this article, I will present an overview of a number of eye-tracking experiments that have evolved from my attempt to re-unite theoretical linguistics and psycholinguistics. The experiments were specifically designed to evaluate the psychological reality of the Primitives of Binding (henceforth *POB*) framework (Reuland 2001), which presents a comprehensive approach to how syntactic structure constrains the use and resolution of anaphoric expressions.

1.1 Theoretical background

Within the linguistic field the restrictions of sentential structure on the use of different anaphoric expressions has traditionally been described through binding theory (Chomsky 1981). Chomsky distinguished three principles: Principle A defines constraints on reflexives and reciprocals (e.g. *himself*, *herself*, *each other*), Principle B constrains pronoun resolution (e.g. *he*, *she*, *his*, *her*) and Principle C constrains full referring expressions and proper names (e.g. *the clown*, *Peter*). One of the crucial implications of Principle A and B is that reflexives and pronouns should appear in a complementary distribution, or put differently, a reflexive should not appear where a pronoun can appear. Although the latter seems to hold in many canonical sentential structures (see (1)), there are some notable exceptions as exemplified in (2)–(4) in which both a reflexive and pronominal yield a grammatically acceptable structure.

- (1) Peter_i hurt himself_i/*him_i.
- (2) Peter_i saw a gun near himself_i/him_i.
- (3) Peter_i said that the queen invited Mary and himself_i/him_i to tea.
- (4) John_i said that there was a picture of himself_i/him_i in the post office.

The numerous violations of the complementarity principle in natural language pose a real problem for binding theory in its classical form and have led other researchers to re-examine Chomsky's claims. As observed by Reinhart and Reuland (1993) (see also Pollard & Sag 1992, 1994 a.o.), the critical notion is whether the anaphoric element and antecedent are arguments of the same predicate. If so, only a reflexive can be used to express the anaphoric relation. Consequently, in Example (1) the occurrence of a pronominal is ungrammatical, since both the anaphoric element and the antecedent are arguments of the predicate *hurt*. However, if the critical elements are arguments of different predicates, reflexives and pronouns are often interchangeable. This is, for example, the case in (1), as the antecedent *Peter* is an argument of the verb *saw*, while

the anaphoric element (*himself/him*) is an argument of the preposition *near*. Reflexives used in a manner as in (2) are referred to as *logophors*.

Reinhart and Reuland proposed that a crucial distinction exists between the interpretation of a reflexive as in (1), in which the antecedent and reflexive are *coarguments* of the same predicate, and reflexives as in (2)–(4), in which they are used logophorically. For the former, the anaphoric dependency can be interpreted on syntactic grounds alone. However, for the latter, extra-syntactic (i.e. discourse related) information is required to determine the proper referent. Hence, in their proposal, Reinhart and Reuland in fact argue for a modular approach to the interpretation of anaphoric dependencies, that is, depending on the exact nature of an anaphoric dependency either the syntactic or discourse module is accessed for the interpretation of that dependency.

Building on the latter proposal, Reuland (2001) presented the outline for the POB framework. He argued that the module responsible for structural computations should be divided into a purely syntactic component and a semantic component. Hence, the architecture of the POB model consists of a syntactic, semantic and discourse module. Similar to Reinhart and Reuland's proposal, the syntactic module is specialized for establishing dependencies between 'true' coargument reflexives and their antecedents by means of A-Chain formation. Without going into further detail, an A-Chain can only be established if the anaphoric element is referentially deficient, which means that it is not fully specified for morphosyntactic features like person, number and gender.¹ For the interpretation of pronominals like *he* and *she* (and logophoric reflexives) on the other hand, two algorithms are available. They are either bound in the semantic module or co-valued in the discourse module. The assumption that the human language system has two independent algorithms for establishing reference between pronominals and their antecedents originated from the ideas of Heim (1982) and is based on two crucial observations (cf. Reinhart 1983; Grodzinsky & Reinhart 1993). One is that quantified (non-referential) antecedents like *every clown* can be equated with a pronoun if the former c-commands the latter, as in (5), but not if such a structural relation is absent, as in (6).

(5) Every clown_i thinks that he_i is funny.

(6) *Every clown_i is very happy. The bearded woman loves him_i.

However, this does not hold if the quantified antecedent *every clown* is substituted with *the clown* (Examples (7) and (8)), which is testament to the claim that different mechanisms are at work in Examples (5) and (8).

1. A more formal definition on A-Chains is: (α, β) form a Chain if (i) β 's features have been (deleted by and) recovered from α , and (ii) (α, β) meets standard conditions on Chains such as uniformity, c-command, and locality (see Reuland 2001, for extensive discussion).

- (7) The clown_i thinks that he_i is funny.
 (8) The clown_i is very happy. The bearded woman loves him_i.

Moreover, the dual approach to pronoun resolution presents a straightforward explanation for the anaphoric ambiguity that arises in sentences with elided verb phrases like (9).

- (9) The acrobat_i likes his_i jokes and the clown_j does $\langle e \rangle$ too.
 (10) The acrobat_i likes his_i jokes and the clown_j does (*likes his_{i,j} jokes*) too.

For present purposes a sufficient approximation of the reconstruction process that underlies the interpretation of these so-called VP-ellipses is to assume a copy-and-paste operation, resulting in a mental representation resembling Example (10). The actual interpretation crucially depends on whether a reader or listener prefers a copy of the VP in which the pronoun *his* is either bound in the semantic module or, alternatively, co-valued in the discourse module. The former results in a reading in which the clown likes his own jokes (*sloppy* interpretation), whereas the latter entails that the clown likes the jokes of the acrobat (*strict* interpretation).

Thus far, I have briefly discussed the differences between classical binding and the POB model at a rather formal level and I suggested that a modular system covers the relevant data better than a more or less unified account. In fact, as the POB model explains the relatively complex pattern of the use of anaphora across a number of languages – in addition to English, Dutch, Icelandic and Frisian also largely conform to its constraints (see Reuland 2001 for discussion) – we may conclude that, from a purely theoretical point of view, the POB framework seems more valid than Chomsky's binding approach. There is, however, another fundamental difference between both approaches, which gives the POB framework an advantage. Namely, whereas the binding theory has no direct bearing on the time-course of the construction of anaphoric dependencies, predictions for real time comprehension follow naturally from the POB's inherent sequential architecture and a general economy principle that governs the division of labor between the different modules. In the next subsection, I will present a more detailed picture of the nature and purpose of the economy principle and, furthermore, address the implications for real time language processing.

1.2 The economy principle

The essence of the POB model can be described in terms of *modularity*, *sequentiality* and *economy*. Modularity refers to the hypothesis that the syntactic, semantic and discourse route should be regarded as autonomous mechanisms, i.e. they are domain specific and functionally encapsulated (cf. Fodor 1983). Furthermore, the POB model also assumes an inherent sequentiality: syntactic operations precede semantic operations,

which themselves are initiated before discourse computations. In this article, however, I will focus on the third core feature, economy. Within the POB framework, economy reflects the idea that the processing resources needed to establish syntactic, semantic and discourse anaphoric dependencies, increase as a function of how many cross-modular steps have to be made. This results in a hierarchy in which syntactic dependencies are more economic than semantic dependencies and the latter in turn require less resources than discourse dependencies (for extensive discussion see Reuland 2001, or Koornneef 2008). In a way, the economy hierarchy (i.e. syntax < semantics < discourse) forms the heart of the POB model. On the one hand, it offers a notational device for explaining the rather complex grammaticality patterns of anaphora across a number of languages and, on the other hand, it allows for specific predictions in terms of real time processing.

To start with its formal function, the economy principle could be interpreted as a linguistic blocking device, similar to the classical Principles A and B, to explain for instance why pronominals cannot be used in a sentence like (1) (i.e. *Peter_i defended himself_i/*him_i*), yet are grammatical in sentences like (5) (i.e. *Every clown_i thinks that he_i is funny*). Its explanatory power follows from the hypothesis that the language system always opts for the most economic anaphoric dependency available and, furthermore, that the more expensive dependencies – in terms of the number of cross-modular steps – are blocked. To illustrate this, in (1) the cheapest way in which the parser can establish reference is via A-Chain formation in the syntactic module. Consequently, the more expensive semantic and discourse route are blocked, which explains why the use of a pronominal like *him* is prohibited. In sentence (5), on the other hand, A-Chain formation is not possible and the semantic route presents the cheapest way for establishing reference, which explains why the use of a pronominal yields a grammatical structure.

In this article, however, I am particularly interested in the actual costs of establishing dependencies within the different modules. More specifically, if the economy principle is not only relevant from a linguistic point of view as a powerful explanatory device, but in addition psychologically real, it predicts the following for real time language processing:

1. The construction of syntactic anaphoric dependencies requires less effort than the construction of extra-syntactic anaphoric dependencies.
2. The construction of semantic anaphoric dependencies requires less effort than the construction of discourse anaphoric dependencies.
3. If a pronominal is ambiguous between a bound (i.e. semantic) and coreferential (i.e. discourse) interpretation, the former is (initially) preferred, since it reflects the most economical option.

I will zoom in on these predictions by presenting the findings of three eye-tracking experiments.

2. Reflexives versus logophors

2.1 Prediction 1

Examples (2)–(4) illustrate that, even though reflexives cannot be replaced by a pronominal in general, there are specifically defined occasions where reflexives and pronominals are interchangeable. A fairly accurate generalization is that the use of both types of anaphoric elements is only allowed if the dependent element and its antecedent are arguments of different predicates and, hence, do not constitute coarguments. According to the POB model, only coargument reflexives are resolved syntactically through A-Chain formation. Logophors, on the other hand, behave like pronominals and are resolved extra-syntactically, either through variable binding or coreference. This fundamental difference between true reflexives and logophors leads us to a very straightforward prediction. Namely, according to the economy hierarchy of the POB model, reflexives should impose fewer costs on the comprehension system than logophors (cf. Prediction 1).

A few cross modal interference studies examined this issue and provided consistent results (Piñango et al. 2001; Burkhardt 2005; Burkhardt et al. 2008). In this dual task paradigm, the participant performs a comprehension and lexical decision task at the same time. Piñango et al. for example, auditorily presented sentences like (11) and (12) for comprehension, which comprised the primary task for the participant. In sentence (11) the reflexive *herself* is a coargument reflexive and in sentence (12) *herself* is used logophorically.

- (11) The girl_i who was arrogant praised herself_i because the network had called about negotiation for a leading role.
- (12) The girl_i sprayed bug repellent around herself_i because there were many mosquitoes in the Everglades.

The secondary task consisted of a lexical decision for a visually presented probe which was entirely unrelated to the sentence. The rationale behind this particular design is that the two tasks compete for the same processing resources. As a result, the reaction time for the lexical decision is an indicator of the processing resources required for the primary task of sentence comprehension. In the experiment, the probe was presented at two critical points in the sentence: at a control position just before the reflexive and at an experimental position immediately after the reflexive. Piñango et al. reported no difference in reaction times at the control position, yet the reaction times at the experimental position were longer in the case of logophors (i.e. in (12)). These

results suggest that the processing of coargument reflexives is more economic than the processing of logophors. More recently, Burkhardt (2005) replicated these results for the Dutch reflexive *zich* (also referred to as *SE* or *simplex* reflexive). Like English *himself/herself*, Dutch *zich* can either be used as a coargument reflexive or alternatively as a logophoric reflexive, and, like Piñango et al. Burkhardt reported longer reaction times in the logophor condition at the experimental position and no differences in reaction times at the control condition. Hence, together the findings from English and Dutch provide preliminary evidence in favor of the first prediction of the regulating economy principle.

Although the results above are in line with the predictions of the POB model, I should address a possible concern regarding the cross modal interference methodology. Its value for measuring processing costs during online language comprehension notwithstanding, the method is somewhat unnatural, as, in day-to-day life people do not normally perform lexical decision tasks while understanding language.² Of course, every experimental setting comes with its limitations. However, Event Related Potential (ERP) studies and eye-tracking studies, for example, may present better, or at least more direct tools for tapping into the language comprehension system in real time. Using the former methodology, Burkhardt (2005, see also Harris et al. 2000) confirmed that in English logophors require more processing resources than coargument reflexives. In order to extend the evidence to Dutch *zich*, I opted for eye-tracking and examined whether the processing differences between logophors and coargument reflexives also emerge in unrestrained reading.

2.2 Experiment 1

2.2.1 Stimuli

Partly based on the items of Burkhardt's cross modal interference study with the Dutch reflexive *zich*, experimental pairs were constructed as exemplified in (13) and (14) below.

(13) Coargument reflexive

De astronaut_i die op Mars een Amerikaanse vlag plantte, verbaasde zich_i toen plotseling een marsmannetje met een nieuwsgierige blik in zijn ogen kwam aanlopen.

The astronaut_i who on Mars an American flag planted, amazed SE_i when suddenly a Martian with a curious look in his eyes came walking by

'The astronaut who planted an American flag on Mars, was amazed when suddenly a Martian with a curious look in his eyes approached.'

2. Note, however, that Piñango et al. (2001), Burkhardt (2005) and Burkhardt et al. (2008) opted for the cross modal interference paradigm, because they specifically assume that in this paradigm sentence comprehension takes place in a normal fashion (Avrutin, personal communication, see also Piñango et al. 1999).

(14) Logophor

De astronaut_i plantte op Mars een grote Amerikaanse vlag naast zich_i toen plotseling een marsmannetje met een nieuwsgierige blik in zijn ogen kwam aanlopen.

The astronaut_i planted on Mars a big American flag next-to SE_i when suddenly a Martian with a curious look in his eyes came walking-by

‘The astronaut_i was planting a big American flag next to himself_i on Mars when suddenly a Martian with a curious look in his eyes approached.’

Like Burkhardt, I included a relative clause in the coargument reflexive condition (see ex. (13)), which enabled me to control for the linear distance between the antecedent and *zich* (i.e. the linear distance in the coargument condition was the same as, or longer than, the linear distance in the logophor condition). Second, the critical verbs were matched for frequency (e.g. in the examples, the verbs *planten* (*plant*) and *verbazen* (*amaze*) were matched, see Burkhardt 2005 for details). Finally, I closely matched the words in the relative clause of the coargument reflexive to the words that appeared before the logophor. From the anaphoric element *zich* onwards the two conditions contained exactly the same material. In order to approach a natural reading setting as much as possible, the stimuli in this experiment (and in Experiment 2 and 3) were presented in their entirety.

2.2.2 Results

If the economy hierarchy for anaphora resolution is reflected in real time language processing, the eye-tracking results should mirror the findings of the dual task and ERP studies. That is, relative to a coargument reflexive, a logophor should cause a delay in reading time right at the anaphoric element – or due to spillover, rapidly thereafter (the term *spillover* is generally used to refer to the phenomenon that the processing consequences of a critical word or region become visible only one or a couple of words downstream in the sentence). Inspection of the eye-tracking data showed a significant first-pass reading difference right at the critical anaphoric element *zich* (t_1 (35) = 2.138, $p = .040$; t_2 (23) = 2.075, $p = .049$).³ As illustrated in Figure 1, *zich* elicited longer reading times in the logophoric condition and, hence, the first prediction was borne out: syntactic

3. Due to size restrictions on the present article, first fixation, first gaze and regression path durations will collectively be referred to as first-pass measures. These first-pass reading times will be contrasted occasionally with second-pass reading times. The latter measure leaves out the initial reading times of a word or region, and as such is informative about any processes that are somewhat delayed in relation to the first encounter of the critical region (cf. Sturt 2003). Second-pass reading times may, for example, be indicative of processes of reanalysis and repair (see Boland 2004, for extensive discussion on how eye movements could be linked to cognitive processes).

anaphoric dependencies impose a higher processing load on the comprehension system than extra-syntactic anaphoric dependencies.⁴




Figure 1. First-pass reading time (in ms) for reflexive and logophor in Experiment 1

3. Variable binding versus coreference

In the two subsequent eye-tracking experiments, we will shift our attention from reflexives to pronominals. Linguistic accounts of pronoun resolution argue that the language system has two ways by which a pronoun is connected to an antecedent; (i) the pronoun behaves as a variable and is bound by its antecedent, or (ii) the pronoun receives a value from the discourse storage through coreference. Variable binding is only licensed in structures with a c-commanding antecedent. Informally, a c-commanding antecedent could be characterized as a sister of the constituent that contains the pronominal.⁵ As can be seen in Figure 2, the noun phrase (NP) *the clown* is a sister of the verb phrase (VP) *thinks that he is funny* because both are dominated by the same mother node, which is in this case the highest sentence (S) node. Hence, in

4. Error bars in Figure 1 and subsequent figures indicate a 95% *inferential* confidence interval, see Tryon (2001).

5. A more formal definition of c-command is as follows: phrase α c-commands phrase β if and only if phrase α does not contain phrase β and the first branching node dominating phrase α also dominates phrase β (see Reinhart 1983, for discussion).

this example the antecedent *the clown* c-commands the pronoun *he* and variable binding is therefore licensed.




Figure 2. Simplified syntactic tree of *The clown thinks that he is funny*

Alternatively, in example (15) and (16) no c-command relation exists, because the antecedent and pronoun appear in different sentences (ex. (15)), or because the antecedent is part of a relative clause and, hence, its position in the syntactic tree is too low to enter into a c-command configuration (ex. (16)). Consequently, in the structures below only a coreferential dependency can be constructed.

- (15) The clown_i is very happy. The bearded woman loves him_i.
 (16) The bearded woman, who is in love with the clown_i, thinks that he_i is very funny.

The second and third prediction of the economy hierarchy pertain to this distinction (i.e. the distinction between variable binding and coreference). First, the focus will be on the second prediction, which was as follows: *The construction of semantic anaphoric dependencies requires less effort than the construction of discourse anaphoric dependencies*. In other words, variable binding should be a cheaper process than coreference.

3.1 Prediction 2

A few studies addressed this second implication of the economy hierarchy. Using the same cross-modal interference paradigm as discussed previously, Piñango et al. (2001) presented sentences like (17) and (18).

- (17) Everyone_i hopes that the tenants will pay him_i the rent before the last day of the month.
 (18) The landlord_i hopes that the tenants will pay him_i the rent before the last day of the month.

In (17) the pronoun *him* refers to the quantified antecedent *everyone* and, because the antecedent is quantified, only variable binding is available. In (18) *the landlord* is not quantified and Piñango et al. assume this will result in a coreferential dependency between the pronoun and its antecedent. In the experiment, the lexical decision target was again presented at two points in the sentence: at a control position just before the pronoun and at the critical position immediately after the pronoun. Piñango et al.

reported no difference in reaction times at the control position, yet the reaction times at the experimental position were shorter in the case of bound-variable pronouns (i.e. (17)). Hence, as predicted by the economy hierarchy, the processing of bound-variable dependencies appears to require fewer resources.

However, more recently Burkhardt (2005) failed to replicate the results in a Dutch equivalent of the Piñango et al. study. In fact, she reported quite the opposite: longer reaction times just after pronouns depending on *iedereen* (the Dutch version of *everyone*). An eye-tracking study in English revealed similar results. Although the experiment was not specifically designed to examine the contrast in processing costs between binding and coreference, Carminati et al. (2002) compared the reading times for pronouns referring to binding antecedents like *every British soldier* to the reading times for pronouns referring to coreferential antecedents like *the old British soldier* and reported longer reading times for the former.

In addition to the results of Burkhardt and Carminati et al. showing an opposite pattern than the results of Piñango et al. another concern with respect to the latter study should be addressed. Although Piñango et al. claim to make a comparison between bound and coreferential pronominals, this is actually not necessarily the case. Recall from our previous discussion that variable binding is always licensed in structures with a c-commanding antecedent. This implies that variable binding is also allowed in sentences like (18) in which the antecedent *the landlord* c-commands the pronominal *him* and, consequently, it remains uncertain whether they are comparing variable binding to coreference or two bound dependencies instead.⁶ This casts some doubt on whether their results truly reflect the processing difference between the two mechanisms or, alternatively, are a consequence of other factors.

In order to overcome this potential concern, Koornneef et al. (2006) conducted an eye-tracking experiment in Dutch in which they compared the reading times for bound pronouns to the reading times for coreferential pronouns. Unlike the previous studies, however, they made sure that the antecedent in the coreferential condition was not in a c-command configuration with the pronoun and, consequently, it is clear that they compared bound to coreferential pronouns. Like Burkhardt and Carminati et al., they reported longer reading times if the pronoun referred to quantified c-commanding antecedents like *every worker* than to a coreferential antecedent like *Paul*. With these results and the remarks I made about the Piñango et al. study in mind, the evidence that suggests that binding requires less resources seems rather limited. Furthermore, we are left with the puzzle of why quantified antecedents induce a relatively small processing load in some experiments, whereas they appear to be relatively expensive in others.

6. The crucial difference between Pinangö et al.'s approach and my approach is that Pinangö et al. contest the existence of generalized quantifiers (Avrutin, personal communication). Hence, in their approach antecedents like *Paul* or *the worker* cannot be bound.

With respect to the latter issue, a possible explanation has been put forward by Burkhardt (2005). She distinguishes between three types of antecedents, namely 'light' quantifiers (e.g. *everyone*), referential quantified antecedents (e.g. *every worker*) and referential antecedents (e.g. *the landlord* or *Paul*). Furthermore, the processing resources that are required to establish a dependency with these antecedents are directly related to what kind of (and how much) information needs to be transferred between the pronoun and the antecedent. A light quantifier like *everyone* requires the smallest amount of processing resources, because only information pertaining to the general set needs to be transferred. A referential antecedent like *the landlord* or *Paul* is a bit more expensive, because it requires the transfer of specific information, which, by hypothesis, is more costly than the transfer of general (variable-like) information. The most expensive dependency is established in the case of referential quantifiers like *every worker*. The crucial difference between light and referential quantifiers is that, while both denote a set of individuals, only the latter provides specific information about the set, through its nominal restrictor (e.g. *worker* in *every worker*). The construction of an anaphoric dependency between a pronoun and referential quantifier thus requires the transfer of both general set and specific information. As a consequence, a referential quantifier is a relatively expensive antecedent, not only compared to light quantifiers but to plain referential antecedents like *the landlord* as well.

Interestingly, most of the experimental results seem consistent with this ranking, with one exception. That is, Burkhardt's Dutch replication study of the Piñango et al. experiment showed an inconsistent pattern, since in her study pronouns referring to the Dutch light quantifier *iedereen* interfered more with the secondary lexical decision task than pronouns referring to plain referential antecedents like *de bakker* (*the baker*). Burkhardt explains this inconsistency by suggesting that *iedereen* and *everyone* do not share the same morphological features and, consequently, *iedereen* should in fact be categorized as a referential quantifier.

Even though the rather puzzling pattern can be explained by adopting Burkhardt's complexity ranking for antecedents (with the inclusion of a stipulation for Dutch *iedereen*), I should raise two concerns regarding the studies conducted so far. First of all, the primary goal of these studies was not to evaluate Burkhardt's antecedent ranking. In fact, matters are the other way around. The ranking is merely used as a post hoc explanation for the reported observations, some of which were unexpected. Moreover, the studies so far only considered two antecedent categories within one experimental design and it may therefore be useful, perhaps even required, to compare all three antecedent categories within one single experiment.

I already addressed the second concern. Especially given the design of Piñango et al. and Burkhardt, it seems impossible to disentangle the effects of antecedent complexity on the one hand, and the choice between binding and coreference on the other. Therefore, it remains unclear whether, for instance, *the landlord* is a more costly antecedent for

a pronoun than *everyone* because (i) they differ in internal complexity, or because (ii) *everyone* binds the pronoun while *the landlord* can only establish a coreferential dependency.

3.2 Experiment 2

3.2.1 Stimuli

In a second eye-tracking experiment, these two issues were addressed by comparing all three antecedent categories within a single design. Furthermore, I made sure that I would be able to attribute the possible effects to either a difference in internal complexity of the antecedent or, alternatively, to a processing advantage for the binding route. Examples of the stimuli are given in (19) to (22) together with their approximate translations.

(19) Light quantified antecedent (Condition 1)

Het was oorlog in Soedan en de soldaten aan de frontlinie waren constant met de dood bezig. Iedereen_i was ontzettend bang dat hij_i zou sterven op het bloedige slagveld.

'A war was going on in Sudan and the soldiers at the front were constantly thinking about death. Everyone_i was very afraid that he_i was going to die on the bloody battlefield.'

(20) Referential quantified antecedent (Condition 2)

Het was oorlog in Soedan en de soldaten aan de frontlinie waren constant met de dood bezig. Iedere soldaat_i was ontzettend bang dat hij_i zou sterven op het bloedige slagveld.

'A war was going on in Sudan and the soldiers at the front were constantly thinking about death. Every soldier_i was very afraid that he_i was going to die on the bloody battlefield.'

(21) Referential c-commanding antecedent (Condition 3)

Het was oorlog in Soedan en een soldaat aan de frontlinie was constant met de dood bezig. De soldaat_i was ontzettend bang dat hij_i zou sterven op het bloedige slagveld.

'A war was going on in Sudan and a soldier at the front was constantly thinking about death. The soldier_i was very afraid that he_i was going to die on the bloody battlefield.'

(22) Referential non-c-commanding antecedent (Condition 4)

Het was oorlog in Soedan en een soldaat aan de frontlinie was constant met de dood bezig. De soldaat_i was ontzettend bang. Hij_i zou sterven op het bloedige slagveld.

'A war was going on in Sudan and a soldier at the front was constantly thinking about death. The soldier_i was very afraid. He_i was going to die on the bloody battlefield.'

The first sentence introduced a group (ex. (19) and (20)) or single character (ex. (21) and (22)). This sentence was primarily included to create more interesting scenarios. In the first three conditions, the second sentence contained the critical anaphoric dependency. The antecedent c-commanded the critical pronoun and was either a light quantifier (i.e. *iedereen* – *everyone*), a referential quantified antecedent (i.e. *iedere soldaat* – *every soldier*) or a plain referential antecedent (i.e. *de soldaat* – *the soldier*). The fourth condition closely matched the third condition, with the crucial difference that the referential antecedent and pronoun appeared in separate sentences. As a result, the binding route is blocked and the parser can only establish a coreferential dependency.

3.2.2 Results

Repeated Measures ANOVAs (Huynh-Feldt corrected) with the four-level factor *Condition* were performed to compare the four experimental conditions (the pairwise post-hoc analyses of the main effects were Bonferroni adjusted for multiple comparisons). According to Burkhardt's complexity ranking for antecedents, pronouns referring to light quantified antecedents like *everyone* should require less processing resources than pronouns referring to plain referential antecedents (e.g. *the soldier*). Similarly, pronouns referring to referential antecedents should require less processing resources than pronouns referring to referential quantified antecedents (e.g. *every soldier*). However, since the experiment was conducted with Dutch sentences and participants, another pattern could arise as well. That is, Burkhardt suggested that *iedereen* should be analyzed as a referential quantifier. Consequently, we may expect that pronouns referring to Dutch light and referential quantifiers are read equally fast, or at least, both should elicit longer reading times than pronouns referring to plain referential antecedents.

None of these possible patterns emerged from the eye-tracking data. In fact, if we only consider the three c-command conditions, the various eye-tracking measures revealed no reliable differences at all. Although the POB model does not make strong predictions for these three conditions, the results are easily explained if the crucial characteristic is not the form of the antecedent, but rather its syntactic relation with the pronoun. In that case we expect the pronoun to be bound in all three c-command conditions and, hence, similar processing costs should emerge.

The reason for including a fourth condition, in which the referential antecedent and pronoun were not in a c-command configuration, was to enable me to tease apart the effects of antecedent complexity and the possibility (or impossibility) of binding the pronoun. By comparing Condition 4 to the minimally different condition in which the same antecedent c-commanded the pronoun (i.e. Condition 3), I could exclusively focus on the difference between binding and coreference. The Repeated Measures ANOVAs for the first-pass eye-tracking measures returned a significant main effect at the two words immediately following the critical pronoun (i.e. the spillover region;

$F_1(2.9, 91.1) = 4.241$, $MSE = 9737$, $p = .008$; $F_2(3, 69) = 6.893$, $MSE = 4615$, $p = .000$). As illustrated in Figure 3, post-hoc analysis revealed that the reading times in the spillover region were longer in Condition 4 than in Condition 3 ($p_1 = .027$, $p_2 = .000$). In other words, binding seems to represent a more economical process than coreference.




Figure 3. First-pass reading time (in ms) for spillover region of Condition 3 and 4 in Experiment 2

Furthermore, the results also suggest that it may be hard to maintain the position that referential antecedents like *the soldier* are always identified through coreference and never through binding (as assumed by Piñango et al. and Burkhardt). This assumption would predict no contrast between pronouns referring to referential antecedents in a binding position and pronouns referring to referential antecedents in a non-binding position. As such a contrast emerged, the data cast some doubt on the validity of the studies conducted by Piñango et al. and Burkhardt, or at least on how they tried to address the predictions of the economy hierarchy.

However, the observant reader might point out that the significant contrast in the spillover region is not due to processing differences between variable binding and coreference per se, but may be related to the sentence boundary in Condition 4. More specifically, it is not inconceivable that the reading time difference in the spillover region reflects spillover from end-of-sentence wrap-up effects and is therefore essentially unrelated to the resolution process of the critical pronoun. According to Rayner et al. (2000), however, it is unlikely that wrap-up effects carry over into the next clause or sentence. In fact, wrap-up seems to be designed to ensure that the information from the sentence is fully integrated, and all within-sentence comprehension problems are

settled before the reader moves on. This provides a way to clear the workspace memory and allows more new information to be brought in. Nevertheless, the alternative explanation based on sentence wrap-up effects *does* illustrate that Condition 3 and 4 are not perfect minimal pairs. Hence, the conclusion that binding represents a more economical process than coreference can only be made tentatively, and similarly, the criticism against Pinangō et al. and Burkhardt's approach should be interpreted with care.

In all, with respect to Burkhardt's complexity ranking for antecedents, we are left with quite a puzzle. On the one hand, a number of studies provided consistent findings, while the present experiment revealed no single result consistent with the hypothesis. At present, I have no straightforward explanation for these conflicting results. However, a safe conclusion is that, although antecedent complexity may play an important role in the course of pronominal resolution, Burkhardt's account may be too strict. Perhaps a more liberal approach is warranted in which antecedent complexity is not defined in terms of a fixed ranking based predominantly on quantification, but instead in more general terms of (semantic and discourse) complexity. Needless to say, further research is required to narrow down this very broad characterization of antecedent complexity.

3.3 Prediction 3

A third prediction that naturally follows from the POB's economy hierarchy is that in case of ambiguity the human parser should initially prefer a bound dependency over a coreferential dependency. Offline studies have provided convincing empirical support for the existence of such a binding preference in different populations. Dutch, English and Chinese children, for example, seem to prefer a bound-variable (or sloppy) interpretation to a coreference (or strict) interpretation and, interestingly, the youngest children exhibit the strongest preference (Foley et al. 1997; Guo et al. 1996; Vasic 2006). The proposed processing hierarchy is also reflected in the performance of agrammatic aphasics with ambiguous ellipsis constructions. Although these patients seemed to be capable of establishing a bound-variable interpretation, they have severe problems with obtaining a coreference interpretation (Vasic 2006). Also for unimpaired adults the strict reading seemed to be the most difficult interpretation to generate offline (e.g. Fiengo & May 1994; Shapiro et al. 2003). However, probably due to having a greater processing capacity they are able to compute both derivations rather quickly and, as a result, their bias towards a sloppy interpretation seems to be somewhat weaker than for young children and agrammatics.

The online evidence, on the other hand, is more limited. To my knowledge, Frazier and Clifton (2000) presented the first online study aimed at examining the hypothesis that variable binding dependencies are preferred over coreferential dependencies. In

a series of self-paced reading experiments, they acquired reading times for the second conjunct of biased ellipses like (23) and (24).

- (23) Sally happened to strain her back yesterday and Fred did $\langle e \rangle$ too. (*sloppy bias*)
- (24) John thinks it's a good idea to shave his face before he goes to sleep and Alice does $\langle e \rangle$ too. (*strict bias*)

In (23) the interpretation is biased towards a sloppy reading (i.e. variable binding interpretation) as it is implausible that Fred strained Sally's back. Alternatively, (24) is biased towards a strict reading (i.e. coreferential interpretation) because it seems unlikely that Alice (a woman) shaves her face before going to bed. They reported longer reading times for the second conjunct in sentences biased towards a strict reading, which suggests that the construction of a coreference dependency requires more processing resources than the construction of a variable binding dependency.

However, they also reported a series of offline questionnaire experiments suggesting that the binding preference of the parser does not generalize to other (quantificational) contexts. For example, sentences like (25) present a similar ambiguity as ellipses because it can either mean *Only Alfred thinks that Alfred is a good cook*, the strict interpretation, or *The only person who thinks of himself as a good cook is Alfred*, the sloppy interpretation.

- (25) Only Alfred thinks that he is a good cook.

If the preference of the parser for binding dependencies is a general one, we should expect the participants to choose the second option more often than the first, but this did not seem to be the case as participants in fact highly preferred the strict interpretation. These, as well as some other results, led Frazier and Clifton to the conclusion that the general binding preference hypothesis "must be abandoned for normal adults" (p.137) as it only emerges within VP-ellipses.

3.4 Experiment 3

3.4.1 Stimuli

Even though Koornneef et al. (2006) showed that the preference for variable binding dependencies is not necessarily restricted to VP-ellipses, but seems to generalize to other contexts as well, Frazier and Clifton's offline results for *only*-structures, which explicitly go against a general binding preference, should be taken seriously. Therefore, I conducted a third eye-tracking experiment in which I presented short stories with a critical third sentence that was ambiguous between a sloppy and a strict reading as illustrated in examples (26) to (29).

(26) Sloppy-bias, only-operator

Lisa en Anouk zijn dol op de muziekkzender MTV. Zij konden hun geluk niet op toen zij mee mochten doen aan het programma 'Pimp My Room', waarin hun kamers werden opgeknapt. Maar helaas, alleen Lisa vindt dat haar gepimpte kamer klasse heeft. Smaken verschillen nu eenmaal.

'Lisa and Anouk love the music channel MTV. They were very happy when they were selected for the show 'Pimp My Room', in which their rooms were redecorated. Sadly, only Lisa thinks that her pimped room has a touch of class. Oh well, each to his own taste.'

(27) Strict-bias, only-operator

Lisa en Anouk zijn dol op de muziekkzender MTV. Lisa kon haar geluk niet op toen zij mee mocht doen aan het programma 'Pimp My Room', waarin haar kamer werd opgeknapt. Maar helaas, alleen Lisa vindt dat haar gepimpte kamer klasse heeft. Smaken verschillen nu eenmaal.

'Lisa and Anouk love the music channel MTV. Lisa was very happy when she was selected for the show 'Pimp My Room', in which her room was redecorated. Sadly, only Lisa thinks that her pimped room has a touch of class. Oh well, each to his own taste.'

(28) Sloppy-bias, ellipsis

Lisa en Anouk zijn dol op de muziekkzender MTV. Zij konden hun geluk niet op toen zij mee mochten doen aan het programma 'Pimp My Room', waarin hun kamers werden opgeknapt. Maar helaas, Lisa vindt dat haar gepimpte kamer klasse heeft, maar Anouk niet. Smaken verschillen nu eenmaal.

'Lisa and Anouk love the music channel MTV. They were very happy when they were selected for the show 'Pimp My Room', in which their roomswere redecorated. Sadly, Lisa thinks that her pimped room has a touch of class, but Anouk does not. Oh well, each to his own taste.'

(29) Strict-bias, ellipsis

Lisa en Anouk zijn dol op de muziekkzender MTV. Lisa kon haar geluk niet op toen zij mee mocht doen aan het programma 'Pimp My Room', waarin haar kamer werd opgeknapt. Maar helaas, Lisa vindt dat haar gepimpte kamer klasse heeft, maar Anouk niet. Smaken verschillen nu eenmaal.

'Lisa and Anouk love the music channel MTV. Lisa was very happy when she was selected for the show 'Pimp My Room', in which her room was redecorated. Sadly, Lisa thinks that her pimped room has a touch of class, but Anouk does not. Oh well, each to his own taste.'

The story was either biased towards a sloppy (ex. (26) and (28)) or strict interpretation (ex. (27) and (29)) and, furthermore, the critical sentence contained the *only*-operator (ex. (26) and (27)) or, alternatively, was an ellipsis that closely resembled the meaning of the *only*-sentence (ex. (28) and (29)). The bias of the story was created in the second sentence by giving some information about both story characters or by giving information about only one story character. The idea was that if you provide information about both,

the interpretation of the ambiguous third sentence is biased towards a sloppy interpretation, but if you provide more or less the same information for only one character the story is biased towards the strict interpretation. An offline pretest of the materials confirmed this assumption.

3.4.2 Results

The eye-tracking data revealed two interesting results. First of all, in the ellipsis region (absent in the *only*-conditions) first-pass eye-movement measures revealed a clear advantage for sloppy interpretations (t_1 (31) = 2.806, p = .009; t_2 (23) = 3.830, p = .001; see Figure 4). These results replicate the findings of the self-paced reading experiments reported by Frazier and Clifton. Hence, across methodologies and languages there seems to be online evidence that readers prefer to assign a sloppy identity to ambiguous elliptic structures.

Furthermore, the reading times for the second sentence (i.e. the sentence that contained the biasing information) also revealed a clear contrast between the strict and sloppy biased stories (F_1 (1, 31) = 39.97, MSE = 94756, p = .000; F_2 (1, 23) = 10.92, MSE = 27658, p = .003). The second-pass durations were much longer (almost twice as long) for the strict biased stories. Interestingly, this was true for both ellipses (t_1 (31) = 4.599, p = .000; t_2 (23) = 2.679, p = .013) and *only*-structures (t_1 (31) = 3.851, p = .001; t_2 (23) = 2.345, p = .028) which suggests that the preference for sloppy interpretations is not restricted to ellipses, but a general property of the parser (see Figure 5). The POB approach can easily account for the striking contrast between sloppy and strict biased stories. Regardless of the context, the human parser should initially prefer a sloppy reading, since binding reflects the cheaper option in the processing hierarchy. Consequently, if the larger context forces a strict reading instead, perceivers have to re-analyze the story to change their initial sloppy reading into the more suitable strict reading. The crucial biasing information in my stories is given in the second sentence, which explains why readers return to that part of the story to resolve the mental inconsistency.




Figure 4. First-pass reading time (in ms) for ellipsis region in Experiment 3




Figure 5. Second-pass reading time (in ms) for second sentence in Experiment 3

4. General conclusion and beyond

4.1 Overview of the results

The economy principle forms the heart of the POB model. On the one hand, it offers a tool to explain the sometimes rather complex grammaticality pattern of anaphoric dependencies across languages. On the other hand, it also makes clear predictions for real time anaphora resolution processes. As discussed, the economy hierarchy is based on the number of cross-modular steps required to assign reference to an anaphoric element. More specifically, for a syntactic dependency less cross-modular steps need to be made than for a semantic dependency, which in turn requires less cross-modular processing than a discourse dependency. The main hypothesis throughout this article was that the economy hierarchy is psychologically real. As a result, one should be able to pick up the difference in processing costs for syntactic, semantic and discourse dependencies in a range of behavioral techniques such as dual task, self-paced reading and eye-tracking paradigms. I embraced the latter methodology and put it to use by testing the three main predictions that naturally follow from the economy hierarchy.

The first prediction was that syntactic anaphoric dependencies should impose fewer processing costs on the comprehension system than extra-syntactic anaphoric dependencies. In Experiment 1 this prediction was tested. The processing costs of the Dutch anaphoric element *zich* were measured in two different structures. In one condition *zich* represented a coargument reflexive and could therefore only be resolved at the syntactic level. In the other condition *zich* was a logophor and could only be interpreted

at an extra-syntactically level. Consistent with the prediction of the POB model, reading times were longer for a logophor than for a true reflexive.

The second prediction of the economy hierarchy implies that the construction of bound dependencies should require less effort than the construction of coreferential dependencies. We have seen that the findings of previous studies were mixed and, furthermore, that some studies reporting consistent results were in fact unable to attribute their results unequivocally to a difference in processing cost for binding and coreference (e.g. Piñango et al. 2001). Bearing this in mind, I conducted a second eye-tracking experiment which allowed me to focus more exclusively on the differences in processing load for the binding and coreference algorithm. Consistent with the POB model, the results of Experiment 2 showed a processing advantage for the former. However, because the relevant experimental conditions were not perfect minimal pairs, I could only conclude tentatively that the second prediction of the economy hierarchy is also reflected in real time measures of language comprehension.

According to the third prediction, readers should have a general preference to bind pronominals in structures that are ambiguous between a bound and a coreferential interpretation. In Experiment 3 this prediction was borne out: readers initially preferred the bound (i.e. sloppy) reading of ambiguous ellipses and *only*-structures, even if the preceding discourse clearly suggested a coreferential (i.e. strict) reading. Moreover, since previous studies revealed similar results, in both offline (Foley et al. 1997; Guo et al. 1996; Vasic 2006; Fiengo & 1994; Shapiro et al. 2003) and online self-paced reading and eye-tracking paradigms (Frazier & Clifton 2000; Koornneef et al. 2006), it seems well-established that in case of ambiguity people have a general tendency to initially choose the semantic route.

4.2 General conclusion

In all, the findings of most studies, including my recent eye-tracking experiments, seem consistent with the three economy predictions of the POB model. Although it is beyond the scope of the present article to fully elaborate on the other two core features of the POB model, sequentiality and modularity, I do note that they are reflected in online measures of anaphora resolution processes as well. An eye-tracking study (Sturt 2003), and dual task studies with healthy adults and agrammatic aphasics (e.g. Nicol & Swinney 1989; Burkhardt 2005; Burkhardt et al. 2008), strongly suggests that coargument reflexives are resolved before discourse information can affect the ongoing comprehension process, i.e. consistent with a syntax-first approach like the POB model. In addition, ERP studies revealed that coargument reflexives and logophors are resolved by qualitatively different mechanisms, a syntactic mechanism in case of the former and a discourse or semantic mechanism in case of the latter (Harris et al. 2000; Burkhardt 2005). In other words, the functional (and temporal) dissociation between

syntactic and extra-syntactic anaphoric resolution mechanisms is possibly also reflected at the neural level: the human brain designates different cortical networks to deal with these different types of dependencies.

The main question posed in the beginning of this article was: Does the POB model have psychological reality? On the basis of my eye-tracking results, as well as the data of other studies – ranging from behavioral measures of the healthy, developing and damaged brain, to neuro-imaging measures – my answer is: I believe it does. It seems therefore legitimate to consider the POB framework as an online account of anaphora resolution. As discussed more thoroughly elsewhere (Koornneef 2008), the most straightforward implementation involves a serial architecture, consisting of three distinct processing phases. In Phase I, the syntactic phase, the processor computes local dependencies in terms of A-Chain formation. This means that reflexives like English *himself* and Dutch *zich* and *zichzelf*, for instance, are linked to their antecedent, but only if the two constitute coarguments of one and the same predicate. If not, the reflexive may still be interpretable, but since these so-called logophoric reflexives behave more or less like pronominals, they depend on variable binding or coreference for their identity instead. The latter two algorithms are executed in the semantic and discourse phase respectively. More specifically, in Phase II of a POB processing cycle, the semantic module may bind pronominals (or logophors), as long as the anaphoric element is in a c-command configuration with its antecedent. If this structural relation is absent, the human parser will be forced to rely on coreference in order to establish an anaphoric dependency. I should note, however, that even though the semantic phase precedes the discourse phase, the eventual choice between a bound-variable or coreferential interpretation of a c-commanded pronoun is intrinsically free, as indicated by the fact that people can easily obtain both a sloppy (i.e. bound) and a strict (i.e. coreferential) reading for VP-ellipses, for instance.

4.3 The next step

This POB-based proposal on real time anaphora resolution becomes particular interesting in the light of a syntax-first neurocognitive model presented by Friederici and colleagues (Friederici 2002; Friederici & Kotz 2003; Grodzinsky & Friederici 2006). Building on the findings of a range of neuro-imaging studies (e.g. ERP, MEG, PET and fMRI studies), they present the outline of a general parsing system (i.e. not specific to anaphora resolution) consisting of three functionally and temporally separable phases of processing. In the first phase, the parser builds a rudimentary hierarchical tree-structure based on purely (lexical) syntactic information, most notably word category information (i.e. noun, verb, determiner etc.). During the second phase, lexical-semantic and thematic role assignment takes place, and a processing cycle ends with a third phase in which the system is engaged in higher level processes such as (syntactic) revision and final integration. Thus, not only the first phase of the

POB and Friederici's model show some remarkable parallels, but, in addition, the subsequent phases are comparable in the sense that more semantic-based and higher level computations are performed.

Given that the POB and the neurocognitive model show a remarkable resemblance to one another, the intriguing opportunity of connecting the functional POB model to the human brain presents itself. More specifically, the neurocognitive model may function as a starting point to address Reuland's (2003) correspondence hypothesis in which "differences between major modules of the grammatical system correspond with differences in processes at the neural level and vice versa" (p.4). Interestingly, I am not alone in my attempt to connect (formal) linguistic operations to Friederici's neurocognitive model. A similar but independently motivated view has recently been put forward by Grodzinsky and Friederici (2006). Partly based on concepts of the minimalist program (Chomsky 1995), they distinguish a number of very basic operations to establish relationships among words and phrases in a sentence: MERGE, MOVE and BIND. Even though Grodzinsky and Friederici do not always explicitly attach these elementary linguistic operations to the three processing phases and associated brain areas of the neurocognitive model, they do seem to sketch the following picture. First of all, they suggest that Phase I is most closely associated with MERGE, which enables the build-up of local phrase structures. During this first syntactic phase, the frontal operculum and anterior superior temporal gyrus (STG) are thought to play a vital role – and in addition the inferior tip of Brodmann's area (BA) 44 may be involved (cf. Friederici 2002). Subsequently, in Phase II BIND and MOVE are executed to establish (anaphoric) dependencies to figure out who is doing what to whom. During this phase Broca's area (BA 44 and 45) appears to be most important. Finally, in Phase III the posterior STG and posterior superior temporal sulcus (STS) are concerned with higher level integration processes, meaning the previously gathered syntactic and semantic information may now interact.

Although their framework resembles the online POB model in some important aspects, most notably in the idea that pronominals are bound in the second phase, their view differs as to when coargument reflexives are resolved. That is, whereas they use the same basic operation BIND to connect both reflexives and pronominals to their antecedents in the second phase of a processing cycle, the POB framework assumes that coargument reflexives are resolved in the first phase through A-Chain formation (which is essentially a MERGE operation). Furthermore, their model does not address how and when coreferential dependencies are resolved and therefore cannot account for the diversity of anaphoric dependencies as found in natural languages.⁷ Nevertheless, bearing these differences in mind, the parallels between their approach and the

7. This should not be taken as criticism, because it was never the intention of Grodzinsky and Friederici to present a full account of anaphora resolution.

POB model are evident, which reinforces the idea that we have the opportunity to formulate some specific hypotheses in relation to where inside the brain the three POB algorithms reside. For instance, since A-Chain formation is essentially a MERGE operation it seems plausible to assume that areas such as the frontal operculum, anterior STG and the inferior part of BA 44 are involved in its execution. Variable binding, on the other hand, more likely requires the involvement of Broca's area, the most important area of Phase II.⁸ For coreference computations matters are less clear-cut, simply because this algorithm makes use of information originating from a lot of different sources (e.g. knowledge about the wider discourse, world knowledge) and consequently it may activate a more diffuse network. In other words, I am hesitant to map coreference only onto the posterior STG and STS – the areas associated with Phase III of the neurocognitive model – particularly because a recent fMRI study has shown that referentially (ambiguous) pronouns recruit medial prefrontal regions as well (Nieuwland et al. 2007).

The next step is then to examine whether the different modules in the POB model indeed resemble the different neural pathways as proposed in the neurocognitive model. Comparing the resolution of true reflexives to logophors, and the resolution of bound to coreferential pronominals in a brain-imaging study, could shed some light on this issue. In fact, the results of these experiments can be used bi-directionally. On the one hand, the POB's modular approach to anaphoric processing is supported if physically distinct networks underlying syntactic, semantic and discourse anaphoric computations can be distinguished. On the other hand, the same results can be used to evaluate the neural circuits as proposed by Friederici and colleagues. Hence, we then create the opportunity to cross-validate two comprehensive models that are very similar in many aspects, even though they evolved in quite distinct disciplines.

In conclusion, I have illustrated how a real time version of the linguistic POB model may meet other models that are currently emerging in psycholinguistics. Obviously, I do not wish to claim that the alliance, or even better, the fusion of the theoretical linguistic and psycholinguistic world will be unproblematic hereafter. All I am saying is that we currently seem to have the theoretical and technological tools to harmonize the advances made in both. Furthermore, the discussion also revealed that in addition to theoretical validity and psychological reality yet another vital hurdle must be taken while dissecting the specifics of human language. Namely, since one of the ultimate goals of language research is – or at least should be – to connect human

8. However, note that Grodzinsky and Friederici suggest that BIND modulates the middle frontal gyrus of the right hemisphere, the middle temporal gyrus of the left hemisphere and the left orbital gyrus.

language comprehension (and production) to the brain, an essential part of any theory is – or should be – its neurocognitive reality.

References

- Boland, J.E. 2004. Linking eye movements to sentence comprehension in reading and listening. In *The On-Line Study of Sentence Comprehension: Eyetracking, ERPs and Beyond*, M. Carreiras & Jr.C. Clifton (eds), 51–76. New York NY: Psychology Press.
- Burkhardt, P. 2005. *The Syntax–Discourse Interface: Representing and Interpreting Dependency*. Amsterdam: John Benjamins.
- Burkhardt, P., Avrutin, S., Piñango, M.M. & Ruigendijk, E. 2008. Slower-than-normal syntactic processing in agrammatic Broca's aphasia: Evidence from Dutch. *Journal of Neurolinguistics* 21: 120–137.
- Carminati, M.N., Frazier, L. & Rayner, K. 2002. Bound variables and c-command. *Journal of Semantics* 19: 1–34.
- Chomsky, N. 1981. *Lectures on Government and Binding*. Dordrecht: Foris.
- Chomsky, N. 1995. *The Minimalist Program*. Cambridge MA: The MIT Press.
- Fiengo, R. & May, R. 1994. *Indices and Identity*. Cambridge MA: The MIT Press.
- Fodor, J.A. 1983. *Modularity of Mind: An Essay on Faculty Psychology*. Cambridge MA: The MIT Press.
- Foley, C., Nunez del Prado, Z., Barbier, I. & Lust, B. 1997. Operator variable binding in the initial state: An argument from VP ellipsis. *Papers on Language Acquisition: Cornell Working Papers in Linguistics* 15: 1–19.
- Frazier, L. & Clifton, C. 2000. On bound variable interpretations: The LF-only hypothesis. *Journal of Psycholinguistic Research* 29: 125–139.
- Friederici, A.D. 2002. Towards a neural basis of auditory sentence processing. *Trends in Cognitive Sciences* 6: 78–84.
- Friederici, A.D. & Kotz, S.A. 2003. The brain basis of syntactic processes: Functional imaging and lesion studies. *Neuroimage* 20: S8–S17.
- Grodzinsky, Y. & Friederici, A.D. 2006. Neuroimaging of syntax and syntactic processing. *Current Opinion in Neurobiology* 16: 240–246.
- Grodzinsky, Y. & Reinhart, T. 1993. The innateness of binding and coreference. *Linguistic Inquiry* 24: 69–101.
- Guo, F., Foley, C., Chien, Y.-C., Chiang, C.-P. & Lust, B. 1996. Operator-variable binding in the initial state: A cross-linguistic study of VP-ellipsis structures in Chinese and English. *Cahiers de Linguistique - Asie Orientale* 25: 3–34.
- Harris, T., Wexler, K. & Holcomb, P.J. 2000. An electrophysiological investigation of binding and coreference. *Brain and Language* 75: 313–346.
- Heim, I. 1982. *The Semantics of Definite and Indefinite Noun Phrases*. Ph.D. dissertation, University of Massachusetts, Amherst. (Published in 1989 by Garland, New York).
- Koornneef, A.W. 2008. *Eye-catching Anaphora*. Utrecht: LOT International Dissertation Series.
- Koornneef, A.W., Wijnen, F. & Reuland, E. 2006. Towards a modular approach to anaphor resolution. In *ESSLI Ambiguity in Anaphora Workshop Proceedings*, R. Artstein & M. Poesio (eds), 65–72.
- Nicol, J. & Swinney, D. 1989. The role of structure in coreference assignment during sentence processing. *Journal of Psycholinguistic Research* 18: 5–19.

- Nieuwland, M.S., Petersson, K.M. & Van Berkum, J.J.A. 2007. On sense and reference: Examining the functional neuroanatomy of referential processing. *NeuroImage* 37: 993–1004.
- Piñango, M.M., Burkhardt, P., Brun, D. & Avrutin, S. 2001. The psychological reality of the syntax-discourse interface: The case of pronominals. In *From Sentence Processing to Discourse Interpretation: Crossing the Borders*, 15–16. Utrecht: Utrecht University.
- Piñango, M.M., Zurif, E. & Jackendoff, R. 1999. Real-time processing implications of enriched composition at the syntax-semantics interface. *Journal of Psycholinguistic Research* 28: 395–414.
- Pollard, C. & Sag, I. 1992. Anaphors in English and the scope of binding theory. *Linguistic Inquiry* 23: 261–303.
- Pollard, C. & Sag, I. 1994. *Head-driven Sentence Structure Grammar*. Chicago IL: CSLI/University of Chicago Press.
- Rayner, K., Kambe, G. & Duffy, S.A. 2000. The effect of clause wrap-up on eye movements during reading. *The Quarterly Journal of Experimental Psychology* 53: 1061–1080.
- Reinhart, T. 1983. *Anaphora and Semantic Interpretation*. London: Croom Helm.
- Reinhart, T. & Reuland, E. 1993. Reflexivity. *Linguistic Inquiry* 24: 657–720.
- Reuland, E. 2001. Primitives of binding. *Linguistic Inquiry* 32: 439–492.
- Reuland, E. 2003. State-of-the-article. Anaphoric dependencies: A window into the architecture of the language system. *GLOT International* 7: 2–25.
- Runner, J.T., Sussman, R.S. & Tanenhaus, M.K. 2006. Processing reflexives and pronouns in picture noun phrases. *Cognitive Science* 30: 193–241.
- Shapiro, L., Hestvik, A., Lesan L. & Garcia, A.R. 2003. Charting the time-course of VP-ellipsis sentence comprehension: Evidence for an initial and independent structural analysis. *Journal of Memory and Language* 49: 1–19.
- Sturt, P. 2003. The time-course of the application of binding constraints in reference resolution. *Journal of Memory and Language* 48: 542–562.
- Tryon, W.W. 2001. Evaluating statistical difference, equivalence, and indeterminacy using inferential confidence intervals: An integrated alternative method of conducting null hypothesis statistical tests. *Psychological Methods* 6: 371–386.
- Vasić, N. 2006. *Pronoun Comprehension in Agrammatic Aphasia: the Structure and Use of Linguistic Knowledge*. Utrecht: LOT International Dissertation Series.

Incremental discourse processing

How coherence relations influence the resolution of pronouns

Willem M. Mak & Ted J.M. Sanders
Utrecht University

The importance of the discourse level for the study of language and linguistics can hardly be overestimated. The study of text and discourse has become an increasingly important area over the last decades, both in linguistics and in psychology. In this paper we report an experiment which adds to our understanding of a crucial process in discourse interpretation: the way in which the integration of clauses takes place. It shows that the position at which the integration takes place depends on the particular coherence relation that holds between sentences. The results once again emphasize the special status of causal relations in discourse processing and representation.

1. Introduction

People use language to communicate. Sometimes, short messages of one word ('Stop!') or one sentence ('I have nothing to say') suffice to get the message across, but usually language users communicate through discourse. The importance of the discourse level for the study of language and linguistics can hardly be overestimated (Sanders & Sanders 2006). For that reason, the study of text and discourse has become an increasingly important area over the last decades, both in linguistics and in psychology.

Discourse is more than a random set of utterances: it shows coherence. Language users make a discourse representation coherent by constructing coherence relations between sentences, such as causal, additive and contrastive relations. These relations can be expressed by connectives like *because*, *and* and *but*. In a cognitive approach to discourse coherence (Sanders & Spooren 2007), connectives are taken as processing instructions: They guide the reader by providing explicit information on the coherence relation that holds between discourse segments. For example, in (1a) the connective *because* signals that the subordinate clause will provide the cause for the situation described in the main clause.

- (1) a. The boy quarreled with his parents, because he did not get permission to go out.
- b. The boy quarreled with his parents. He did not get permission to go out.

The same causal relation between the clauses can be inferred in (1b), but here the reader has to connect the two clauses without the help of a connective providing information on the meaning of the relation; in such cases readers infer the relation between the discourse segments.

Given the proposed function of connectives as “processing instructions”, it is an important question what the effect is of using a connective between two clauses. Millis and Just (1994) proposed the Reactivation Hypothesis, which claims that connectives evoke integration of the clauses and reactivation of the first clause. Millis & Just argue that this reactivation takes place at the end of the second clause. They conclude this from a series of experiments in which they used probe recognition tasks to assess whether words from the first clause were active during processing of the second clause. They found evidence for this accessibility at the end of the clause, but not in the middle of the clause.

In contrast, Traxler et al. (1997) claim that the influence of connectives is already noticeable as the second clause unfolds. They compared the processing of causal (2a) and diagnostic (2b) sentences.

- (2) a. Heidi felt very proud and happy, because she won first prize at the art show.
- b. Heidi could imagine and create things, because she won first prize at the art show.

Readers had greater difficulty processing diagnostic sentences. This difference in processing difficulty was visible *before* the end of the sentence, at the region directly after the verb (*first prize* in the example). The authors conclude that readers compute the relationship between the clauses incrementally, during the reading of the second clause.

While investigating the processing of causal connectives, Noordman and Vonk (1997, 1998) distinguish between three functions of the connective *because* in the reading process. First, it signals how the material is structured, the *segmentation function*. Second, it signals to the reader that the two sentences are causally related, the *integration function*. Third, it is an instruction to the reader to infer the extra-textual information that underlies the causal relation, the *inference function*.

The effect of these three functions of the connective is studied by Cozijn (2000; see also Cozijn et al. 2003). He compared the processing of texts with and without the connective *because* (sentences 3a and 3b).

- (3) a. John experienced a big delay. There was a traffic jam on the highway.
- b. John experienced a big delay, because there was a traffic jam on the highway.

Compared to the sentence without a connective, Cozijn found a speeding up effect at the region following the connective (*was a traffic jam*) and a slowing down at the sentence-final region (*on the highway*). Cozijn attributed these effects to the integration function and the inference function, respectively. The connective makes it easier to integrate the two sentences, hence the early speeding up effect of the connective. Cozijn claims

that the connective “signals that the sentence consists of a consequence-cause construction and that the words following it should be integrated into the cause-part of the structure” (p. 80). The slowing down at the end of the sentence is attributed to the inference process: The connective elicits the inferring of the major premise of the causal chain of reasoning. According to Cozijn, this inference is not automatically made in the absence of the connective.

Thus, there is conflicting evidence concerning the time course of integration. A second unresolved question is what cognitive processes take place during the integration and the inference stage. We investigated these two issues in an experiment in which we used pronoun resolution as a means of tapping into the processes involved in integrating consecutive sentences. We used sentences that were connected with a connective. The first sentence introduced two protagonists (see sentences 4a–4d). In a pretest, we established that readers had a preference for the protagonist occupying the subject position (Noun 1) to be the subject of the second clause. In the second sentence one of the protagonists is referred to with a pronoun. This pronoun either refers to Noun 1 (see sentence 4a) or to the protagonist occupying the object position (Noun 2, see sentence 4b).

- (4) a. The teenager quarreled with his parents, because he did not get permission to go out.
- b. The teenagers quarreled with their father, because he did not give permission to go out.

In these two sentences, it is immediately clear to whom the pronoun refers. These two sentences were then compared to sentences with the same content, but with a pronoun in the second clause that was ambiguous as to which referent it refers to (4c and 4d).

- (4) c. The teenagers quarreled with their parents, because they did not get permission to go out.
- d. The teenagers quarreled with their parents, because they did not give permission to go out.

It is well known from the literature that readers immediately look for an antecedent of a pronoun (e.g., Koornneef & van Berkum 2006). The question is what readers do with the information on the pronoun to process the relation between the sentences. If this integration process takes place at a late stage (Millis & Just 1994), the Noun 2-completion sentences (4b and 4d) should show processing difficulty irrespective of the disambiguating information of the pronoun. However, if the integration takes place incrementally (Cozijn 2000; Traxler et al. 1997), the readers may use the disambiguating information of the pronoun in 4b to integrate the sentences immediately. The processing difficulty at the end of the sentence of Noun 2-completions may then be absent, or at least be smaller. In the first case, one would predict a main effect of the completion, in the latter case, one would predict an interaction: processing difficulty at the end of the sentences with Noun 2 completion, but only when the pronoun is ambiguous.

For the sentences with an unambiguous pronoun, the effect of N1-preference may already be visible at the pronoun at the beginning of the subordinate clause. In that case, longer reading times should be found at the pronoun in (4b), compared to (4a). However, it may also be the case that an unambiguous pronoun prevents the reader from pursuing the wrong interpretation at all. In that case, there will be no effect of N1-preference at that point, but the N1-preference should only lead to difficulty at or after the verb of sentences like (4d), in which the verb disambiguates towards N2 as the subject of the subordinate clause.

2. The experiment

2.1 Method

Participants. 40 students of the University of Utrecht participated in the study. They were paid for their participation.

Materials. Forty sets of four sentence-pairs were constructed (see Table 1 for an example). The first sentence introduced two entities, one of which was the subject of the sentence, the other the object.

The second clause started with a connective followed by a personal pronoun that referred to either the subject or the object of the first clause. The semantic content of the verb in the second clause always made clear which of the entities was referred to by the pronoun. In addition, the relative pronouns were manipulated in such a way, that the number marking on the personal pronoun either did or did not make it possible to identify the correct antecedent.

Table 1. Example of a set of sentences used in the experiment. The point of disambiguation is given in bold

Condition 1: N1 antecedent, disambiguation at the pronoun.
De journalist volgde de politici, omdat hij de verkiezingen versloeg voor de krant.
<i>The journalist followed the politicians, because he covered the elections for the newspaper.</i>
Condition 2: N1 antecedent, disambiguation at the verb.
De journalisten volgden de politici, omdat ze de verkiezingen versloegen voor de krant.
<i>The journalists followed the politicians, because they covered the elections for the newspaper.</i>
Condition 3: N2 antecedent, disambiguation at the pronoun.
De journalisten volgden de politicus, omdat hij de verkiezingen manipuleerde door omkoping.
<i>The journalists followed the politicians, because he manipulated the elections by bribing.</i>
Condition 4: N2 antecedent, disambiguation at the verb.
De journalisten volgde de politici, omdat ze de verkiezingen manipuleerden door omkoping.
<i>The journalists followed the politicians, because they manipulated the elections by bribing.</i>

As a result, the sentences were either disambiguated by the number marking of the personal pronoun (conditions 1 and 3), or by the semantic content of the verb (conditions 2 and 4).

We performed a pretest, in order to ensure that we had items in which there was a preference for N1 to be the antecedent of the pronoun. Sixty-seven participants received a booklet, in which a main clause was presented, followed by a connective and a pronoun. Both nouns in the first clause were plural, and so was the pronoun. The participants were asked to complete the second clause. We computed the percentage of completions in which the pronoun referred to N1. We selected 40 items for which at least 80% of the completions were N1-completions.

The 40 experimental items and 60 fillers were pseudorandomly divided into 5 blocks of trials. Four experimental versions were constructed. The items occurred in the same order in each version. Across the experimental versions each item occurred in every condition. The participants saw each item only once, and saw 10 experimental items in each condition. At the beginning of the experiment there was a practice block of 10 trials. To make sure that the participants read the sentences carefully, verification statements were included after 25% of experimental and filler items.

Apparatus. The materials were presented on a PC monitor. The eye movements were recorded using an Eyelink I eye tracker. The position of both eyes was sampled with a frequency of 250 Hz.

Procedure. Before the experiment, the participant was informed about the purpose of the apparatus. Then the headset was placed on the participant's head. Each block, including the practice block, started with a calibration. At the beginning of a trial the participants saw a fixation point, indicating where the sentence would begin. When the participants looked at this fixation point, the sentence appeared on the screen. As soon as they had finished reading the sentence, the participants pressed a button. At that moment they either saw the fixation point of the next sentence, or a verification statement about the sentence they just read. When there was a verification statement, the participants indicated their judgment by pressing a button, after which the fixation point of the next sentence appeared.

2.2 Results

Five regions in the target sentence were analyzed: These are illustrated in (5). The first region contained the connective and the pronoun, the second region contained the object of the second clause, the third contained the verb, the fourth contained a preposition and an article, and the fifth contained a noun.

- (5) De journalist volgde de politici,/ omdat hij/ de verkiezingen/ versloeg/ voor de/
krant./

Two measures were computed. The first measure was the *first-pass reading time*. This is the time spent in a region in first-pass, before the region is left with a progressive saccade (going forward to a later region) or with a regressive saccade (going back to an earlier region). The second measure was the *regression path duration* which consists of the sum of fixation and saccade durations from the time when the reader encounters a region, to the time when the reader enters the region after that region. We analyzed the data using an analysis of variance, with two factors: Noun Attachment (N1 versus N2) and Disambiguation (disambiguation at the pronoun versus disambiguation at the verb).

No effects were found at Regions 1 and 2.

Region 3, Verb: There was a significant main effect of Noun Attachment (first-pass reading times: $F(1,37) = 7.04, p < .05$; $F(1,39) = 2.87, p = .10$; regression path duration: $F(1,37) = 5.99, p < .05$; $F(1,39) = 4.23, p < .05$). The verb was read faster when the personal pronoun referred to the first noun, than when it referred to the second noun.

Table 2. First-pass reading times

<i>Antecedent Disambiguation</i>	N1 at pronoun	N1 at verb	N2 at pronoun	N2 at verb
1. Connective + Pronoun	302	311	301	305
2. Noun Phrase	322	316	324	315
3. Verb	261	279	292	297
4. Preposition + Article	322	330	352	425
5. Final Noun	500	518	503	503

Region 4, position after the Verb: There was a main effect of Noun Attachment (first-pass reading times: $F(1,34) = 20.85, p < .001$; $F(1,38) = 10.85, p < .005$; regression path duration: $F(1,34) = 16.92, p < .001$; $F(1,39) = 7.10, p < .05$) and of Disambiguation (first-pass reading times: $F(1,34) = 12.93, p < .001$; $F(1,38) = 3.54, p = .07$; regression path duration: $F(1,34) = 10.61, p < .005$; $F(1,39) = 8.07, p < .01$).

Importantly, these main effects were qualified by a significant interaction of Noun Attachment and Disambiguation (first-pass reading time: $F(1,34) = 6.90, p < .05$; $F(1,38) = 2.84, p = .10$; regression path duration: $F(1,34) = 4.06, p = .05$; $F(1,39) = 8.04, p < .01$). There was no difference between the conditions in which the pronoun referred to N1 (first-pass reading time: both $F_s < 1$; regression path duration: $F(1,37) = 1.87, p = .18$; $F(1,39) < 1$). When the pronoun referred to N2, reading times in the sentences with disambiguation at the verb were longer than reading times in the sentences with disambiguation at the pronoun (first-pass reading time: $F(1,36) = 16.94, p < .001$; $F(1,38) = 4.99, p < .05$; regression path duration: $F(1,34) = 10.56, p < .005$; $F(1,39) = 13.66, p < .005$).

Region 5, sentence-final Noun: There was a main effect of Disambiguation in the regression path duration: $F(1,34) = 4.33, p < .05$; $F(1,38) = 5.18, p < .05$.

Table 3. Regression Path Duration

<i>Antecedent Disambiguation</i>	N1 at pronoun	N1 at verb	N2 at pronoun	N1 at verb
1. Connective + Pronoun	308	313	301	310
2. Noun Phrase	390	390	395	394
3. Verb	305	305	342	336
4. Preposition + Article	521	585	633	894
5. Final Noun	1502	1714	1647	1747

3. Discussion

Two conclusions seem to follow from the data: First, the main effect of type of Noun Attachment at the verb shows that N2-completions are more difficult than N1-completions overall. Second, the interaction at the position after the verb shows that the difficulty of the N2-completion is modified by the presence of a disambiguating pronoun: In the cases with an ambiguous pronoun the N2-completion led to longer reading times, in the cases with an unambiguous pronoun there was no significant difference between the reading times of N1-completions and N2-completions. These results indicate that an ambiguous pronoun is attached by the reader to N1, which leads to processing difficulty at the verb of N2-completions, whereas an unambiguous pronoun guides the reader towards the correct interpretation.

There were no effects at the pronoun. This may imply that the disambiguation is early enough to override the N1-preference. However, it is also possible that there was some difficulty for N2-completions disambiguated at the pronoun, but that this difficulty was not visible in our eye movement measures.

3.1 Effect of connective and relation type

In previous experiments it has been shown that the integration effect depended on the particular connective that was used, and on the type of relation that holds between the sentences. Mulder (2008) reports an experiment in which materials are used that are similar to the materials used by Cozijn (2000). Mulder compared the effect of the Dutch connectives *omdat* and *want*. Both can be translated to English with “because”, but the connective *want* mostly signals a subjective causal relation between the clauses, whereas *omdat* is used both for subjective and objective causal relations (see, among

others, Pit 2003). Mulder only finds a facilitating effect at the word directly following the connective (*there* in example 3), but does not replicate the speeding up effect at the following region. At the last region of the sentence, he finds the slowing down effect for the connective *want*, but for *omdat* a slowing down effect is found later, in the first region of the next sentence. This suggests that the different connectives lead to different effects. In other experiments, the relation between the sentences is varied. In these experiments, the effect of the causal connective is different for connectives signaling a subjective relation versus connectives signaling an objective relation. This was found for English (Kamalski et al. 2008) and for Dutch (De Leeuw et al. 2008). These experiments show that a speeding-up effect is found when the causal connective signals an objective relation, as in (6a), but not when the causal connective signals a subjective relation, as in (6b).

- (6) a. Er was een aardbeving. Daardoor stortte het gebouw in
There was an earthquake. As a result, the building collapsed.
- b. Er was een aardbeving. Dus zal het er wel gevaarlijk zijn.
There was an earthquake. Therefore it must be dangerous there.

Hence, the way in which a connective influences the integration of the second sentence with the first depends both on the particular connective that is used and on the type of coherence relation there is between the sentences. In our experiment, most of the sentences contained the connective *toen* (“when”; 28 cases) or the connective *omdat* (“because”; 9 cases). To investigate whether these different connectives had a different effect, we will first look at the connectives *omdat* and *toen* separately. After that we will investigate the difference between sentences that were causally related and sentences that were not causally related.

3.2 Effect of connective

To investigate whether different connectives led to different processing patterns, we analyzed the *omdat*-cases and the *toen*-cases separately. Because of the low number of items in the *omdat* cases, we could not do an analysis by participants, so we only performed an analysis by items. We concentrated on the positions and measures where we found effects in the main analysis, the verb and the two following positions.

At the verb, in the main analysis a main effect of Noun Attachment was found. This effect was completely absent for the *omdat*-cases (first-pass reading times: difference 1 ms; $F_2 < 1$; regression path duration: difference 5 ms, $F_2 < 1$), but it was present for the *toen*-cases (first-pass reading times: difference 25 ms; $F_2(1,27) = 4.94$, $p < .05$; regression path duration: difference 74 ms; $F_2(1,27) = 4.77$, $p < .05$).

Figure 1 shows the results of the position after the verb. In the main analysis an interaction of Noun Attachment and Disambiguation was found. In the *omdat*-cases

this interaction reappeared (first-pass reading times: $F(1,7) = 10.53$, $p < .05$; regression path duration: $F(1,8) = 8.66$, $p < .05$) but for the 28 *toen*-cases there was no interaction (first-pass reading times: $F < 1$; regression path duration: $F(1,27) = 3.00$, $p = .095$). There was only a significant main effect of Noun Attachment (first-pass reading times: $F(1,27) = 11.59$, $p < .005$; regression path duration: $F(1,27) = 5.45$, $p < .05$).

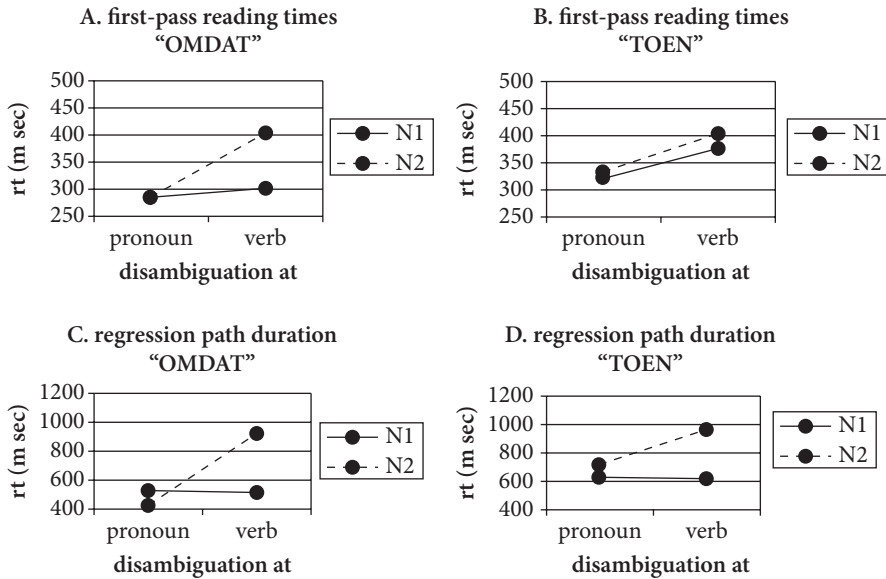


Figure 1. Results at the position after the verb, for the connectives *toen* and *omdat* separately

At the sentence final position, in the main analysis an effect of Ambiguity was found. This effect was not significant for the *omdat*-cases nor for the *toen*-cases separately.

The simplest explanation for the difference in the results for the *omdat*-cases versus the *toen*-cases is that readers process these two connectives differently. However, although the interaction that is present in the overall analysis is not significant for *toen*, inspection of the means in the regression path duration (Figure 1D) raises doubts whether the interaction was really absent. It may well be that the interaction was only weaker, and therefore not significant. There is another reason for doubting the result for *toen*. In an ERP study by Brehm (2005), readers read a lead-in sentence like (7a), followed by a target sentence like (7b), that either started with the causal connective *darum* (therefore), the temporal connective *danach* (afterwards), or the sentential adverb *gestern* (yesterday).

- (7) a. Das Auto war auf dem Sandweg steckengeblieben.
 The car had got stuck on the sandy path.
- b. Darum/Danach/Gestern beschaffte Niklas Kies für die Auffahrt.
 Therefore/Afterwards/Yesterday Niklas got gravel for the driveway

Brehm found a difference in processing between the connective conditions on the one hand, and the sentential adverb condition on the other hand. There was no difference in processing between the sentences with a causal connective and the sentences with a temporal connective. This seems to be at odds with the difference between *omdat* and *toen* in the present experiment. However, this difference may not be due to the connective per se, but to the actual relation readers establish between the sentences. Brehm's experimental materials allow for a causal interpretation of the temporal connective *danach*. Even though this is strictly speaking a temporal connective, the readers probably inferred a causal relation between the sentences.

3.3 Effect of relation

On the basis of the discussion above, one may hypothesize that the processing difference between *toen* and *omdat* is not caused by the connective per se, but rather by the interpretation of the connective (Kamalski et al. to appear; De Leeuw et al. submitted). This should not come as a surprise. The explicit linguistic code that we use in everyday language is often underspecified. We know that language users often interpret linguistic items differently in different contexts, and that people systematically make inferences. This also holds for connectives. Spooren (1997) has shown how coherence relations are not always matched by the meaning of the connectives in the context. More specifically, he has shown how language users deal with so-called underspecified coherence relations, in examples like (8).

- (8) Sinds Feyenoord door Opel gesponsord wordt, heeft de ploeg geen wedstrijd meer gewonnen.
Ever since Feyenoord is sponsored by Opel, the team has not won one game.

Here, the temporal connective *sinds* will receive a causal interpretation. This type of inferential mechanism is powerful. In the proper context, connectives like *and*, *when* and *after* regularly receive a more specific interpretation, for instance a causal instead of an additive or temporal one (see Spooren 1997 for a discussion of these issues). Likewise, the connective *toen* can not only be interpreted as expressing a temporal relation, but also as marking a causal relation. Now, if it is indeed the actual interpretation of the connective that causes the difference in processing, then we would expect that the cases of *toen* that allow for a causal interpretation (from now on referred to as "causal") would be processed like the *omdat*-cases, whereas the cases of *toen* that cannot be interpreted causally show a different pattern. We checked this prediction in our materials. Of the 28 cases of *toen* in the experiment, 14 could be interpreted as marking a causal relation; the other 14 could not get a causal interpretation and were therefore likely to be interpreted temporally. We analyzed the data for these two sets of items separately. At the verb there was no main effect of Noun Attachment for causal

toen (first-pass reading times: difference 20 ms; $F(1,13) = 1.40$, $p = .26$; regression path duration: difference 61 ms, $F(1,13) = 1.02$, $p = .33$), but a significant main effect for non-causal *toen* in the regression path duration (first-pass reading times: difference 31 ms; $F(1,13) = 3.74$, $p = .08$; regression path duration: difference 88 ms, $F(1,13) = 6.80$, $p < .05$).

Figure 2 shows the results for the position after the verb for causal and non-causal *toen* separately. For causal *toen*, the same pattern of results was found as for *omdat*: there was a significant interaction of Noun Attachment and Disambiguation in the regression path duration (first-pass: $F(1,13) = 4.03$, $p = .07$; regression path: $F(1,13) = 4.67$, $p = .05$). There was no difference between the sentences in which the personal pronoun referred to N1 (first-pass reading times: $F(1,13) < 1$; regression path duration: $F(1,13) < 1$). However, when the personal pronoun referred to N2, the sentences disambiguated at the verb led to significantly longer reading times than the sentences disambiguated at the personal pronoun (first-pass reading times: $F(1,13) = 6.48$, $p < .05$; regression path duration: $F(1,13) = 6.18$, $p < .05$).

The pattern of results for non-causal *toen* was different (see Figure 2): There was only a main effect of Noun Attachment (first pass: $F(1,13) = 10.60$, $p < .01$; regression path: $F(1,13) = 5.54$, $p < .05$). The sentences that were disambiguated towards N2 attachment were read slower than the sentences that were disambiguated towards N1 attachment, irrespective of how they were disambiguated.

At the sentence-final position no difference was found between causal and non-causal *toen*.

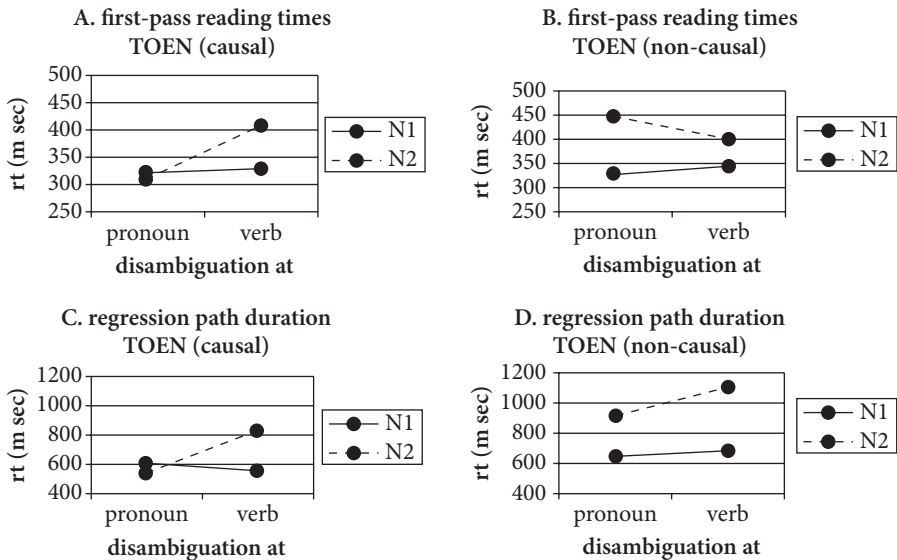


Figure 2. Results at the position after the verb, for causal *toen* and non-causal *toen* separately

This pattern of results shows that it is not the connective per se that causes the sentences to be processed differently, but the preferred interpretation of the connective: the processing of causal *toen* patterns with *omdat*, where non-causal *toen* leads to a different pattern. For the connective *omdat*, only one interpretation is possible: it signals that the two clauses are causally related. The connective *toen* can both be interpreted as signaling a causal relation and as signaling a non-causal relation. The analysis above showed that when *toen* is used as a causal connective, processing of the second clause is similar to the clauses with *omdat*, whereas when *toen* is used as a non-causal connective the processing pattern is different.

4. Summarizing discussion

The results can be summarized as follows: When the sentence is disambiguated at the verb, processing difficulty is found when the sentence has N2 as the subject. This means that readers, on encountering the ambiguous pronoun, attach it to N1. The off-line preference for N1 to be the subject of the second clause is thus visible in the reading time data. When the sentence is disambiguated at the pronoun, the picture becomes more complex: In the case of sentences that allow for a causal relation, the difficulty of continuations with N2 as the subject completely disappears. In the case of sentences that cannot receive a causal interpretation, the difficulty of N2-completions remains the same, in spite of the clear disambiguation at the pronoun. The difference between clauses that are causally related and clauses that are not causally related also becomes clear in the position in the sentence where the effects appear. For the non-causally related clauses, the effect already appears at the verb, whereas the effect appears at the position after the verb for the causally related sentences.

There are two important findings that need further discussion. The first is the absence of any effect in the region containing the pronoun, the second is that the effects at and after the verb are modified by the specific relation between the clauses.

The absence of an effect at the region containing the pronoun is surprising. One might conclude from this that an N2 completion is not problematic at all, if the disambiguating information is given at the beginning of the clause. However, this is inconsistent with the results of Koornneef and Van Berkum (2006). In an experiment on implicit verb causality they found an effect of consistency immediately at the pronoun in sentences like (9a) and (9b). The verb *praise* leads to a bias for N2 to be the subject of a subordinate clause beginning with *because*. In (9a) the pronoun is consistent with this bias, in (9b) it is inconsistent.

- (9) a. Linda praised David because he...
- b. Linda praised David because she...

Koornneef and Van Berkum found longer regression path durations at the pronoun in (9b) compared to the pronoun in (9a). A similar effect is not found in our experiment. A possible reason for this difference is the fact that different relations were processed differently in our experiment. In some cases, the case marking on the pronoun led to the disappearance of the effect at the verb, whereas in other cases the effect at the verb was unaffected by the case-marking on the pronoun. This difference in the effect of the pronoun may have blurred any effects that are present in only a part of the items.

The most important finding in the experiment concerns the results at and after the verb. These results show a difference between clauses that are causally related and clauses that are not causally related. For the clauses with *omdat* and the clauses with causal *toen*, no effects are found at the verb. The only effect that is found is at the position after the verb, and this effect is different from the effect found for non-causal *toen*. Here, the effect of violation of the N1-preference is only found when the pronoun is ambiguous. Our explanation of this result is as follows: When the clauses are connected with the causal connective *omdat*, readers know that there is a consequence-cause relation between the clauses, and that the clause starting with *omdat* gives the cause part of that relation. Hence the reader can immediately integrate the two clauses, without having read the second clause. Apparently, when the clauses are integrated immediately, the disambiguating information of the pronoun is sufficient to override the N1-preference.

In the case of the connective *toen*, the reader does not know what relation there is between the clauses. The relation may just be an additive or a temporal relation, without any further connection between the two. However, the connective *toen* can also be used when there is a causal relation between the clauses. The results show, that if the first clause is likely to be the consequence of an action, readers process the second clause in the same way as the clauses with *omdat*. Apparently, if a consequence-cause relation is likely, readers prefer to integrate the clauses in this way.

If the first clause does not call for a cause, the process is different. As the reader does not know what kind of relation will follow, the integration of the clauses can only take place when the main proposition of the second clause is known. In our sentences this is the case when readers process the verb. Integration is thus done at that point, and this integration is difficult if the subject of the second clause is different from the subject of the first clause. Since the integration is made this late, the case-marking on the pronoun does not override the N1-preference.

An important point to note here, is that the difference in processing that we found is not due to the connective *per se*, but is apparently due to the coherence relation readers infer between the sentences. More specifically, readers seem to interpret this relation in a causal way, when the first segment allows them to. This may explain why in the ERP-experiment by Brehm (2005) no difference in processing was found between sentences connected with a causal connective and sentences connected with a temporal

connective. In her experiment, Brehm used the same sentences, and only changed the connective between the sentences. As was the case in our “causal *toen*”-items, these sentences could be causally related (otherwise the causal connective could not have been used) and hence the absence of a result is in line with the results of the present experiment.

What is the most important finding in our experiment? We presented ambiguous pronouns in the second discourse segment, to gain insight in on-line discourse processing: These pronouns could refer to N1 or N2 in the first segment. We first established that in this case the processing of N2-completions was indeed problematic. However, in the case that the pronoun was unambiguous, we found that the type of relation between the sentences influenced online processing. When readers expected a causal relation – i.e., when they process a *because* or a causally interpretable *when*-relation – they did not suffer from any difficulty when N2 was the subject of the second clause. By contrast, *when*-relations that could not be interpreted causally *did* raise processing problems, even if the linguistic information of the pronoun was sufficient to resolve it successfully. We conclude from this that when readers expect a causal relation between sentences, they integrate them before the proposition of the second clause is known.

This account allows for several important conclusions. First, it shows that the coherence relation that readers infer between sentences immediately influences the interpretation of the second segment that is connected by this relation. Similar findings are reported by Kehler (2002), who following on the seminal work by Hobbs (1979), has shown how different coherence relations give rise to different ways of pronoun resolution. Second, it suggests that readers tend to interpret these relations causally whenever the first segment allows them to. Both conclusions fit in remarkably well with a cognitive account of coherence relations (Sanders 2005; Sanders & Spooren 2007).

Two issues remain open: First, our account implies that the expectation that readers have about the relation of the clauses guides their interpretation of the connective *toen*. The question is whether the presence of the connective is crucial for this effect. One may hypothesize that the same processes will occur when there is no connective between the clauses. We cannot exclude that possibility on the basis of the present experiment. This issue can be resolved by doing an experiment in which the same sentences are used, but without the connective. If the connective is crucial for the integration process to occur immediately, the difference between causally and non-causally related causes should disappear.

A second issue is whether the causal connective *omdat* is sufficient for the early integration of the clauses. In that case, early integration should be found even when the first clause is not recognizable as a consequence.

The results from our experiment underline the special status of causal relations in discourse processing. Language users seem to have a preference for connecting

information causally (Noordman & Vonk 1998; Sanders & Noordman 2000) because their reading strategy is to construct a highly connected representation. Indeed, when confronted with a sequence like *Bill entered the room. Bob left*, it is tempting to infer a *Cause-Consequence* relation: *Bill entered the room. That's why Bob left*. This preference also seems to hold in the reverse order: *Bob left. Bill entered the room* is interpreted as *Bob left because Bill entered the room*. Formal discourse representation approaches have suggested a principle of *Maximize Discourse Coherence*, according to which the 'rich' causal (*Explanation*) relation is to be preferred over the additive one (Asher & Lascarides 1998: 107). In a similar vein, Levinson (2000: 122) when discussing conjunction, notes that: "when events are conjoined, they tend to be read as temporally successive, and if at all plausible, as causally linked". He does not present this as an on-line processing preference, but these arguments might give rise to an alternative hypothesis, which we call the *Causality-by-default hypothesis* (Sanders 2005): Because (experienced) readers aim at building the most informative representation, they start out assuming the relation between two consecutive sentences is a causal relation (given certain characteristics of two discourse segments). Subsequently, causally related information will be processed faster, because the reader will only arrive at an additive relation if no causal relation can be established. We are currently in the process of further refinement and testing of this tempting hypothesis (see also Mulder 2008, Chapter 5).

A preference to connect sentences in a causal way would explain the fact that readers process the information from the causal *toen*-cases in the same way as the *omdat*-cases, even though there is no explicit marker of the causal relation between the clauses. Note, however, that if the first clause is not recognized as a consequence, readers apparently do not assume that the subordinate clause will be causally related to the main clause.

In conclusion, the results of the present experiment add to our understanding of a crucial process in discourse interpretation: the way in which the integration of clauses takes place. It shows that the position at which the integration takes place depends on the particular coherence relation that holds between sentences. The results once again emphasize the special status of causal relations in discourse processing and representation.

References

- Asher, N. & Lascarides, A. 1998. Bridging. *Journal of Semantics* 15: 83–113.
- Brehm, E.U. 2005. Connective ties in discourse: Three ERP-studies on causal, temporal, and concessive ties and their influence on language processing. Ph.D. dissertation, Potsdam University.
- Cozijn, R. 2000. Integration and inference in understanding causal sentences. Ph.D. dissertation, Tilburg University.

- Cozijn, R., Vonk, W. & Noordman, L.G.M. 2003. Afleidingen uit oogbewegingen: de invloed van het connectief 'omdat' op het maken van causale inferenties. *Gramma/TTT* 9: 141–156.
- Hobbs, J.R. 1979. Coherence and Coreference. *Cognitive Science* 3:67–90.
- Kamalski, J., Lentz, L., Sanders, T. & Zwaan, R.A. 2008. The forewarning effect of coherence markers in persuasive discourse: evidence from persuasion and processing. *Discourse Processes* 45: 323–345.
- Kehler, A. 2002. *Coherence, Reference, and the Theory of Grammar*. USA: CSLI Publications.
- Koornneef, A.W. & Van Berkum, J.J.A. 2006. On the use of verb-based implicit causality in sentence comprehension: Evidence from self-paced reading and eye tracking. *Journal of Memory and Language* 54: 445–465.
- Leeuw, S. de, Mak, P. & Sanders, T. 2008. *Effects of the Dutch Causal Connectives 'Dus' and 'Daardoor' on processing*. Toegepaste Taalwetenschap in Artikelen: 53–63.
- Levinson, S.C. 2000. *Presumptive Meanings. The Theory of Generalized Conversational Implicature*. Cambridge, MA: MIT Press.
- Millis, K.K. & Just, M.A. 1994. The influence of connectives on sentence comprehension. *Journal of Memory and Language* 33: 128–147.
- Mulder, G. 2008. Understanding causal relations. Ph.D. dissertation, Utrecht University.
- Noordman, L.G.M. & Vonk, W. 1997. The different functions of a conjunction in constructing a representation of the discourse. In *Processing Interclausal Relationships: Studies in the Production and Comprehension of Text*, J. Costermans & M. Fayol (eds), 75–93. Mahwah, NJ: Erlbaum.
- Noordman, L.G.M. & Vonk, W. 1998. Memory based processing in understanding causal information. *Discourse Processes* 26: 191–212.
- Pit, M. 2003. *How to Express Yourself with a Causal Connective Subjectivity and Causal Connectives in Dutch, German and French*. Amsterdam: Rodopi.
- Sanders, T. 2005. Coherence, Causality and Cognitive complexity in Discourse. Paper presented at the First International Symposium on the Exploration and Modelling of Meaning SEM-05, in Biarritz.
- Sanders, T.J.M. & Noordman, L.G.M. 2000. The role of coherence relations and their linguistic markers in text processing. *Discourse Processes* 29: 37–60.
- Sanders, T. & Sanders, J. 2006. Text and Text analysis. In *Encyclopedia of Language and Linguistics, 2nd Edition*, K. Brown et al. (eds), 597–607, London: Elsevier.
- Sanders, T. & Spooren, W. 2007. Discourse and text structure. In *Handbook of Cognitive Linguistics*, D. Geeraerts & J. Cuykens (eds), 37–60. Oxford: Oxford University Press.
- Spooren, W. 1997. The processing of underspecified coherence relations. *Discourse Processes* 24: 149–168.
- Traxler, M.J., Bybee, M.D. & Pickering, M.J. 1997. Processing causal and diagnostic statements in discourse. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 23: 88–101.

Theoretical validity and psychological reality of the grammatical code

Ad Neeleman & Hans van de Koot*
University College London

One interpretation of the competence-performance distinction is that the grammar is a knowledge base consulted by the performance systems. We argue for an alternative interpretation, according to which the grammar and the performance systems are theories of the same object, but at different levels of description. This does not mean that the grammar is *reduced* to the performance systems. It is fundamental in explaining properties of these systems, as it expresses constraints on their design that generalize across specific computational procedures. We base our argument on considerations of computational efficiency, a comparison with other information-processing systems and evolutionary considerations.

1. Ristad's problem

Much progress in theoretical linguistics has been made possible by the distinction between competence and performance, which is essentially an idealization over the data that linguists work with. Chomsky (1965: 3) formulates it as follows:

- (1) Linguistic theory is concerned primarily with an ideal speaker-listener, in a completely homogeneous speech-community, who knows its language perfectly and is unaffected by such grammatically irrelevant conditions as memory limitations, distractions, shifts of attention and interest, and errors (random or characteristic) in applying his knowledge of the language in actual performance.

The idealization proposed here is hardly controversial: in all empirical sciences it is customary to ignore noise in the data.

A competence theory must of course be supplemented by a characterization of its relation to the systems of performance. Chomsky (1965: 9) argues that the relation

*For useful comments we would like to thank two anonymous reviewers, as well as Peter Ackema, Michael Brody, Dirk Bury, Janet Dean Fodor, Vieri Samek-Lodovici, Øystein Nilsen, Colin Phillips, Neil Smith, Michal Starke, Peter Svenonius, Kriszta Szendrői, and the participants of TiLT 2004 and of the 2003 Current Issues in Syntax class at UCL.

between the grammar and the systems of comprehension and production can be rather abstract (see also Newmeyer 2003):

- (2) When we say that a sentence has a certain derivation with respect to a particular generative grammar, we say nothing about how the speaker or hearer might proceed, in some practical or efficient way, to construct such a derivation. These questions belong to the theory of language use – the theory of performance.

This view certainly goes beyond the common practice of ignoring noise in the data in insisting that the theory of competence does not necessarily tell us anything about what happens in comprehension and production. While Chomsky's quote in (2) is nearly forty years old, it is an accurate description of the current state of affairs: there is no transparent relation between minimalist transformational grammar and plausible models of human language computation. For example, language computation is presumably left-to-right, while grammar – if taken to be derivational at all – is cyclic. In conjunction with widely held assumptions about phrase structure, this implies that derivations proceed largely right-to-left.

Although we accept both aspects of the competence-performance distinction, we think that Ristad (1993) is justified in pointing out that this distinction threatens the coherence of the generative enterprise. If competence theory has nothing to say about the computations performed by speakers and hearers, one may well ask what object the theory describes (see especially Ristad 1993: 110–117 for discussion). In what follows we will refer to this question as Ristad's problem:

- (3) *Ristad's problem*
What object is a competence grammar a theory of, given that it abstracts away from language computation?

The main aim of this paper is to compare two answers to this problem. According to the first, the grammar is a knowledge base consulted by the performance systems; according to the second, the grammar is a description of the language faculty at a more abstract level than that at which the performance systems are best described. We will argue in favour of the latter view.

2. Two answers


There are a number of tasks that need to be carried out before a speaker can utter, or write down, a sentence (see Levelt 1989 for discussion and further references; see Dennett 1992 for some critical remarks). We will refer to the system responsible for these tasks as the *generator*; it maps conceptual-intentional information to a set of instructions for the system of articulation (abstracting away from written or signed languages). Similarly, there are various computations that must take place before

someone who hears a sentence can understand what is being said (see Pritchett 1992; Gorrell 1995; and Frazier & Clifton 1995 for discussion and further references). The system responsible for these computations, usually referred to as the *parser*, maps an auditory signal to conceptual-intentional information.

One might argue that the grammar is a separate module of the brain consulted in some fashion by the performance systems. Chomsky (2000: 117) seems to suggest an architecture of this type:

- (4) There is good evidence that the language faculty has at least two different components: a “cognitive system” that stores information in some manner, and performance systems that make use of this information for articulation [and] perception [...].

This quote does not give an explicit characterization of the relation between competence and performance. One model compatible with it is given in (5), where the double arrows are meant to represent the information flow from grammar to parser and generator.



This model provides a straightforward justification for Chomsky’s idealization in the quote in (2) and thereby suggests an answer to Ristad’s problem. Generative linguistics is the study of the knowledge module at the center of (5). Linguistic behavior is the result of the interaction of this module with the generator and parser. Hence, if one wants to find out about the linguistic module, one should abstract away from various aspects of this behavior.

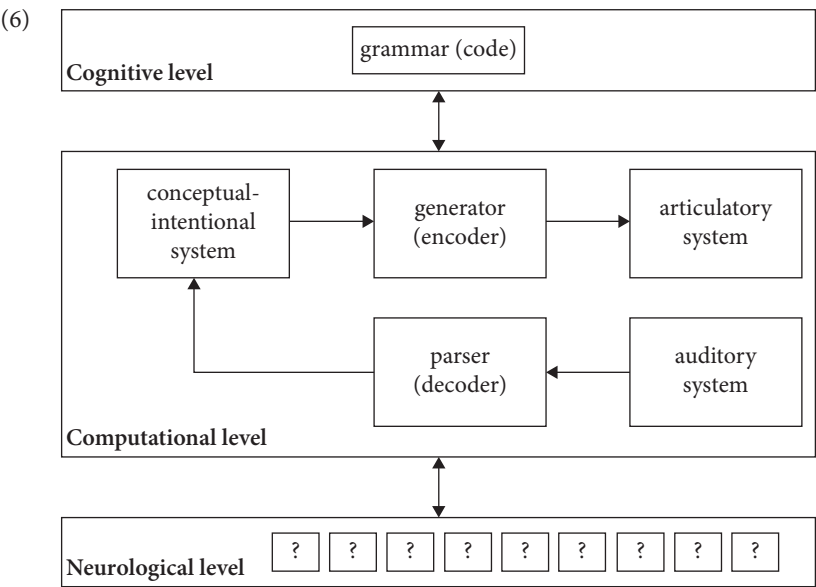
An alternative to the view that the grammar is a knowledge module at the same level of description as the performance systems is suggested by Marr’s (1982) discussion of complex information-processing systems. Marr argues that any machine carrying out an information-processing task must be understood at three levels. The most abstract level is a description of the device in terms of the logical structure of the mapping that it carries out from one type of information to another. In the case of humans, we can refer to this as the cognitive level of description. And in the case of language

it is the natural locus of hypotheses about regularities in the mapping between sound and meaning.

The next level down is a description of the algorithm that yields the desired input-output mapping. We will call this the computational level.¹ Of course, the representations associated by the algorithm must stand in a relation of logical equivalence (or something approaching it) to the descriptions at the cognitive level (see Abney 1988; and Van de Koot 1990 for discussion). But the algorithm itself may bear little resemblance to the mathematical characterization of the mapping at the cognitive level. In the case of language the computational level consists of a description of the performance systems.

At the most basic level, we must offer a description of how the algorithm and its input and output are realized physically. In the case of human language, this is a theory of the neurology associated with linguistic computation.

The following model depicts this tripartite description of the human language faculty. The generator and parser are modules at the computational level, responsible for encoding what a speaker wants to say and for decoding what a hearer receives as input. The grammar is a description at the cognitive level of the logical structure of these computations, a kind of linguistic code.



The main difference between the models in (5) and (6) is that the former assumes that the theories of competence and performance characterize two objects at the same level

1. What we call the cognitive level corresponds to Marr’s computational theory. What we call the computational level corresponds to Marr’s algorithmic level.

of description. This must be the case, because the model crucially assumes information flow between the grammar and the parser/generator. In contrast, the model in (6) assumes that the theories of competence and performance characterize a single object at different levels of description. Both proposals constitute answers to Ristad's problem, but with very different implications.

Those readers only familiar with recent work in transformational grammar – the so-called minimalist program – may feel that the model in (6) is a dramatic departure from generative orthodoxy. In fact the opposite is true. Marr himself points out that Chomsky's theory of transformational grammar belongs to the most abstract level of description. As he puts it (Marr 1982: 28):

- (7) [F]inding algorithms by which Chomsky's theory may be implemented is a completely different endeavor from formulating the theory itself. In our terms, it is a study at a different level, and both tasks have to be done.

One might think that developments since the publication of *Aspects of the Theory of Syntax* have invalidated this view, an impression strengthened by the way in which syntacticians often talk about derivations as actual computations and of the syntax as the computational system. Consider, for instance, the motivation for the notion of 'phase' (comparable to the more traditional notion of cycle), which plays an important role in current minimalist thinking. As Chomsky (2005) makes clear, phases are assumed to limit the search space for linguistic operations such as movement, thereby reducing the load on short-term memory.

Outside the derivational tradition in linguistics, the model in (6) is probably much less controversial. Many psychologists, for example, would agree that the grammar is not a knowledge base consulted in actual language computation. However, those who hold this view often tie it to an instrumentalist interpretation of grammar, according to which grammatical theory is not about anything real, but at best a handy way of talking about the performance systems. In fact, there are also several linguists who have relativized the importance of grammatical theory in this way (see Phillips 1996; Li 1997; Kempson et al. 2001; Culicover & Nowak 2003; among others).

In what follows we will combine a defence of the model in (6) with an argument for a realist view of grammar. Generalizations that hold at the cognitive level of description are of critical importance to understanding linguistic computation (and vice versa). Before elaborating on these issues, we will argue that the model in (5) simply cannot work.

3. On grammar consultation

An evaluation of the model in (5) requires a viable characterization of the consultation relation between grammar and performance systems. Given that the parser and

the generator build linguistic representations, we must ask ourselves what the role of the grammar is in this process. For the idea of consultation to have empirical content, it must be the case that the performance systems are capable of generating structures that the grammar rejects. If the performance systems only build structures that adhere to the grammar, then there is no plausible function left for a separate grammatical module at the same level of description.

The simplest conceptualization compatible with these requirements treats parser and generator as minimally specified structure-builders, possibly enriched with a number of task-specific strategies. The outputs of these structure builders form the input to the grammar, whose principles act as filters that either reject the input as ill-formed or pass it on to language-external systems. This proposal can be summarized as in (8a) below.

There are alternatives in which the notion of grammar consultation has progressively less empirical content. These can be derived from (8a) by interpreting some principles of grammar as part of the structure-building mechanisms, rather than as filters (see (8b)). If all principles are reinterpreted in this way (see (8c)), it is fair to say that grammar consultation is entirely redundant, as parser and generator are unable to build structures that violate the grammar.

- (8) a. All principles of grammar are filters on structures proposed by parser and generator.
- b. Some principles are built into parser and generator, other principles are filters.
- c. All principles are built into parser and generator.

A reasonable interpretation of (8a) assumes a parser that is limited to binary branching structures and that obeys feature cooccurrence restrictions (such as the restriction that no node can belong to two different categories). Such a parser can generate a very large number of structures for even simple inputs. For a three-word input, the number of trees will be equal to $2c^2$, where c is the number of possible feature combinations that can characterize a node. This complexity has two sources, namely tree geometry and the content of nodes. The parser has a choice of two tree shapes (this yields the factor 2). In addition, a minimal tree for a three-word sentence contains two complex nodes, each of which can have c feature combinations (this yields the factor c^2). The number of possible trees grows very rapidly with every additional word, since the rate of growth in the number of tree geometries is almost factorial, while the rate of growth in the number of node content assignments is exponential (given 7 mutually compatible features, c equals 128, which implies that there are $5 \cdot 128^4$ potential trees – around $13 \cdot 10^8$ – for an input string of 4 words).

While the set-up just sketched is a strawman, it allows us to see that a simple generate-and-test design yields so many candidate structures that any computational

system is quickly overwhelmed by it. There are three main strategies for addressing this issue. (These are not mutually exclusive and in practice researchers typically employ more than one of them.) The first is to further constrain the parser, so that fewer candidate structures can be built. The second is to call on filters as early as possible rather than at the end of the structure-building process. Finally, one may be able to limit the number of candidate structures by underspecifying them. We discuss the merits of each of these strategies in turn, limiting ourselves to the parser (all conclusions we draw carry over straightforwardly to the generator).

4. The benefits of grammar compilation


The transfer of principles from grammar to parser can only be beneficial if the computational problems of the original generate-and-test design are not replicated internally to the parser. That is to say, this module should not itself consist of an unconstrained structure assigner plus a set of filters (for discussion of this point in the context of Government and Binding Theory, see Kolb & Thiersch 1991). Therefore, many researchers adopt an approach in which inputs are pre-parsed using a set of context-free rewrite rules derived from various principles of the competence grammar. (Indeed, parsing with such a system requires no more than cubic time; see Younger 1967; and Earley 1970.) The structures recovered in this way are then presented to a set of filters that correspond to other principles of the competence grammar. This strategy of translating grammatical principles into low-level rules is usually referred to as grammar compilation (see Marcus 1980; Berwick & Weinberg 1984; and Abney 1989, 1991).

Grammar compilation has advantageous computational effects because it eliminates the search space associated with very general and abstract constraints by recoding them in a (possibly large) number of locally applicable rules. Let us illustrate this using Inclusiveness, the main principle regulating the content of nodes in current theories of syntax (Chomsky 1995a,b; the formulation below is taken from Neeleman & Van de Koot 2002):

(9) *Inclusiveness*


Properties of a nonterminal node must be recoverable from its daughters; properties of a terminal node must be recoverable from the lexicon.

Inclusiveness allows upward transfer of features, as in (10a), as long as no node is skipped, as happens in (10b). It does not allow downwards or sideways transfer of features, as in (10c) and (10d).




The effects of Inclusiveness can be encoded in a context-free grammar if every feature that is part of the left-hand side of a rewrite rule recurs in the right-hand side of the rule. We illustrate how this can help reduce computational load in parsing using a toy example. Consider a hypothetical language whose grammar can be characterized as follows: (i) trees are uniformly right-branching binary structures; (ii) the feature content of nodes is governed by Inclusiveness; (iii) percolation of features can only take place from the right; (iv) all features must percolate to the root node; (v) a node can contain one of c possible feature combinations.

Consider first a set-up in which the above conditions on the percolation of features are built into the parser in the form of a set of context-free rewrite rules that emulate feature inheritance from right daughters (the general format would be $N_i \rightarrow X N_i$, where N_i is any node with a well-formed feature specification). Given that the number of possible feature combinations is c , the compiled grammar will consist of a maximum of c^2 rewrite rules (namely the number of possible values for X multiplied by the number of possible values for N_i). Suppose we are parsing the string ABC and have not yet encountered C (see the partial structure in (11)). The maximum number of trees that the parser will need to consider upon parsing C is c . Once A has been found, the parser can postulate c alternative partial structures, given that its mother can be the left-hand side of only c rewrite rules. Once B has been found, the parser can expand each of these partial structures in exactly one way, because B 's mother and A 's mother are guaranteed to be identical in feature content (in (11) this is indicated by the index i). Upon encountering C , the parser only needs to determine which of the c possible feature combinations is contained in that terminal. Therefore, irrespective of the length of the input string, the search space for the parser never exceeds c . In conclusion, this very restricted grammar is not associated with any serious computational load, if compiled in the manner outlined above.



Suppose instead that Inclusiveness and right-headedness are filters and that the parser builds right-branching structures only. This arrangement leaves the parser free to postulate any permissible feature combination for any node; illegal patterns of percolation are only filtered out once a complete tree is presented to the grammar. If this parser analyzes the string ABC, the number of trees to be considered upon encountering the final input symbol will be c^2 . This is because nothing in the parser links the feature content of the nodes that immediately dominate A and B (see (12)). In general, the parser's search space at the end of a parse will be c^{n-1} , where n is the number of symbols in the input string. This growth rate is the result of a generate-and-test design in which Inclusiveness does not directly restrict the options of the parser but is only called upon when a complete structure has been built.



The reductions in search space associated with grammar compilation are illustrated in (13), which gives an overview of the number of structures assigned to strings using increasingly constrained parsers.² The table is based on a grammar that is more realistic than the toy example just considered. It allows projection from either the left or the right and a mixture of left- and right-branching structures. Hence, trees can be distinguished on the basis of tree shape, as well as the content of nodes. In order to keep the example workable, we assume that each lexical item has a categorial feature, and an index that uniquely identifies it. Feature percolation is restricted to projection of this information.

Column A shows how many structures the parser allows for an input string if constrained by binary branching and a weak form of Inclusiveness that allows feature percolation to skip nodes, as in (10b), but disallows sideward or downward transfer.

2. These numbers have been calculated using a Prolog program containing a parameterized merge predicate. This program can be downloaded at www.phon.ucl.ac.uk/home/hans/merge.pl.

The results in column B obtain if the assignment of a label to a nonterminal node takes place under direct domination, as required by the strict version of Inclusiveness in (9). Although this represents a marked improvement, the order of growth in column B is still unacceptable, as it outstrips that of 2^n in column D, included here as a benchmark. Adding directionality of label assignment to the system leads to a further decrease in the number of alternative structures (although the growth in column C is still exponential). The dramatic improvement shown in column E is a consequence of introducing selection as a further restriction on structure assignment.

(13)

	A	B	C	D	E
Men (1)	1	1	1	2	1
Men slept (2)	2	2	1	4	1
The man slept (3)	12	8	2	8	1
Men bought a book (4)	112	40	5	16	1
A man bought a book (5)	1360	224	14	32	1
Men said men bought a book (6)	19872	1344	42	64	1
Men said a man bought a book (7)	? ³	8448	132	128	1

A: merge with weak Inclusiveness and binary branching

B: merge with Inclusiveness and binary branching

C: merge with Inclusiveness, binary branching and label assignment under directionality

D: 2^n

E: merge with Inclusiveness, binary branching, label assignment under directionality and selection

The constraints added to the parser's structure-building device in (13) correspond to well-known grammatical restrictions. Label assignment under direct domination yields the merge operation standardly adopted in minimalism (see Chomsky 1995a,b). Directionality of label assignment mimics the effects of directionality restrictions, such as those encoded by word order parameters. Selection, finally, reflects the existence of argument structure.

It is, in fact, widely acknowledged that efficiency in parsing can be improved by grammar compilation (see Berwick & Weinberg 1984; Fong 1991; and the papers in Berwick et al. 1991). If compilation is necessary, the question arises what the status is of compiled principles. Once a principle has been compiled, it need not be present in the grammar as a filter, because no candidate structures can be generated that violate it. This conclusion of course bears on the validity of the model in (5): the more principles are compiled, the less support there is for the view that the grammar is a module

3. The question mark in this cell is due to limitations of the computers available to us.


at the same level of description as the parser. Hence, proponents of (5) must find other ways of making parsing efficient.

5. The benefits of interleaving


The inefficiency of grammar consultation could be alleviated by allowing intermediate structures to be presented to the grammar, so that unwanted structures are eliminated early on. This technique is often referred to as ‘interleaving’. The crucial question in the current context is whether interleaving can be an alternative to compilation. There are circumstances in which this is the case, such as the toy example considered earlier (we will not demonstrate this here). In general, however, interleaving cannot remove all sources of computational intractability.

The key obstacle is that almost all intermediate structures are ungrammatical. For example, a partial structure for a transitive sentence that lacks an object violates the θ -criterion. Hence, this principle cannot be checked until all arguments have been integrated into the tree.

A problem of this type also arises in the evaluation of partial structures for adherence to Inclusiveness. Inclusiveness cannot be checked if the direct domination relations in the right edge of a tree might be altered when the parser extends the structure. A feature F that threatens to violate Inclusiveness in a node N may turn out to have to its origin in an as yet unintegrated terminal or partial tree that may end up as a daughter of N . Hence, there might be a grammatical continuation of the parse, even though the initial tree seems to violate Inclusiveness. (14) illustrates this scenario, where both β and γ might be complex categories.



It follows that Inclusiveness can be applied to a partial structure only if it has a right sister that contains at least one terminal node. This guarantees that the right edge of the partial structure in question, consisting of δ and β in (15), cannot undergo further modification.



We may conclude that, as long as direct domination relations cannot be fixed at intermediate stages of the parse, there will be cases in which Inclusiveness cannot be interleaved with structure-building operations. This gives rise to unacceptable complexity associated with node content, even if interleaving of other principles succeeds in eliminating the complexity associated with tree shape. In particular, no right-branching structure can be evaluated with respect to Inclusiveness unless one of two conditions is met. Either the parser has reached the end of the input string or it hypothesizes that the right-branching structure becomes part of a left-branching structure. But this means that in the worst case the number of candidate trees that are presented to the grammar is $O(c^n)$, where, as before, c is the number of possible feature combinations and n is the number of terminals dominated by the right-branching (sub-)tree. If c equals 128, this implies that upon integration of the final terminal of the string *I like old port*, well over 6 million candidate structures must be considered by Inclusiveness. This example suggests that the beneficial effects of interleaving will be modest.

Even this may be too optimistic, however. As our example demonstrates, intermediate structures often violate some principle or other. Therefore, successful interleaving of grammatical principles with parsing operations requires the existence of auxiliary procedures that determine whether a given partial structure can be presented to a given principle. Fong (1991) shows that the computational load that results from the application of these procedures in many cases outweighs the benefits of interleaving. He concludes, in the context of a GB-based parsing system, “[...] that parsing times will increase dramatically as more principles are interleaved”. The key problem is that the auxiliary procedures do not only apply to partial structures that ultimately give rise to a successful parse, but also to those that eventually fail. This point suggests that interleaving is not a viable alternative to compilation. The number of failing structures that must be considered is larger if the structure builder is less restricted. The structure builder is less restricted if fewer principles have been compiled.


6. The benefits of underspecification

As just discussed, Inclusiveness – if conceived of as a filter – can only apply to a given structure if either the final input symbol is reached or the structure is placed on a left branch. This allows the parser to generate too many candidate structures before the principle can have its filtering effect.

Intuitively, a parser that allows underspecification of node content might be able to circumvent this problem. The idea would be that, rather than building many fully specified trees, the parser initially builds only a single underspecified one. Once enough information is available, node content is fixed in an efficient manner. Such a

strategy can contribute to computational efficiency, but does not constitute an alternative to compilation.

To see what the potential benefits of underspecification amount to, consider a uniformly right-branching structure that is deemed complete and has been put on a left branch, as in (16a). At this point, there are c choices for the content of the node directly dominating B and C . Assuming an interleaved application of Inclusiveness, one of these will be selected (let us say N_i ; $N_i = C$ if the grammar is right-headed). Next, there are c choices for the node that immediately dominates A , of which again one will survive application of Inclusiveness. Therefore, the overall number of trees that the parser presents to Inclusiveness for a right-branching structure dominating n nodes will never exceed $n \cdot c$, which seems acceptable.



The overall picture is, however, less rosy than this narrative suggests. The essence of the underspecification strategy is to avoid commitment to any particular analysis. To begin with, there is abundant evidence that language computation has a high degree of incrementality, probably in order to associate the input with an interpretation as early as possible. This is apparent from garden path phenomena, amongst other things.

Moreover, the discussion of (16) presupposed commitment to tree shape in the absence of commitment to node content. But it is hard to see on what basis a parser could decide on a particular tree shape if all nonterminal nodes in the tree lack features. For example, the constituent containing A , B , and C may be placed on a left branch because it is selected by D . However, that selectional relation requires that the relevant constituent have selectable properties. In other words, its root node cannot remain underspecified. Much the same is true of the internal structure of the constituent under discussion. In other words, if the parser truly fails to commit to the content of nonterminal nodes, it must also fail to commit to tree shape. This implies that, when D is encountered, the state of the parse is that in (17); that is, an unanalyzed string.



A cursory inspection of the literature suffices to establish that those researchers who rely on underspecification do indeed also make use of compilation in order to avoid

postponement of commitment. For example, Marcus 1987; Marcus et al. 1983; Berwick and Weinberg 1985; and Gorrel 1995 all compile at least X' -theory (the precursor of Inclusiveness) and lexical selection (θ -theory and c -selection).

7. The status of the grammar

The discussion so far suggests that a description of the language faculty at the computational level should acknowledge two modules, namely a generator (or encoder) and a parser (or decoder), each employing a compiled special-purpose rule system. This implies a very much reduced role for a separate module of grammar at this level of description: it would be largely redundant. We thus arrive at the model in (6), in which the grammar is a description of the language faculty that abstracts away from computation, a description at what we called the cognitive level. This model allows a reconciliation of results in theoretical and computational linguistics. Grammatical theory is valid as a description of the language code, while special-purpose rule systems are part of the encoding and decoding algorithms (of course, in the model in (6) these special-purpose rule systems cannot be derived by a process of grammar compilation, as the grammar is not a component of the same level of description).

Since both the generator and the parser contain a rule system that instantiates the grammar, one may wonder what the point is of claiming mental reality for grammatical principles in addition to this? Indeed some linguists believe that there is no need to introduce a separate level of description for the grammar. The maxim that the grammar is the parser (endorsed by Phillips 1996, 2003; Kempson et al. 2001; and Steedman 2000; amongst others) is indicative of a research program along these lines. We are sympathetic to this kind of work insofar as it is aimed at developing a model of the language faculty at the computational level. After all, it directly addresses the concerns expressed above. We believe, however, that the principles of competence grammar have a right of existence independently from their implementation in parser and generator.

What is our justification for this claim? In general, a level of description is justified if it captures generalizations that cannot be captured at alternative levels of description. Marr gives as an example the case of some gas in a bottle. If one wants to describe the thermodynamic properties of such a mini system (temperature, pressure and volume), and the relationships between these properties, one does not use a large set of equations that each describe the movements of one of the particles involved. Rather we use a higher level of description to answer questions that could otherwise not be addressed. This is not just for practical reasons. The laws of thermodynamics are statistical in nature and, as a matter of logic, statistical statements cannot be derived from

non-statistical premises, unless a specific theory linking them is formulated (see Popper 1992 for extensive discussion).

In fact, it is fair to say that in certain complex systems the organization of lower levels of description is explained by properties that only reveal themselves at higher levels of description. A simple example is that of human artifacts. These always have a function that, at least to some extent, explains their design and hence their physical realization. But functions cannot easily be described at the level of physics. Similar considerations extend to evolutionary development, where one can talk about the function of an organ, its design (given that function), and its physical instantiation (given that design).

Such top-down explanations are never the whole story, as characteristics of higher levels can often be realized in different ways at lower levels. In the case of artifacts there might be various designs that realize a function and various physical realizations of a design. Which choices are made depends on various extraneous considerations, such as fashion, and the availability and price of certain materials; other relevant factors may be given by limitations at lower levels of description, such as the laws of physics. In the case of evolution the direction of development is partly determined by properties of the organism at earlier stages, as well as by properties of the environment. This notwithstanding we feel that any attempt at explaining artifacts or evolution must partly be top-down, if only because this allows us to distinguish those aspects of physical realization that are crucial from those that are coincidental. An exclusively bottom-up (or reductionist) approach must fail in this respect.

Before we work out a more detailed motivation for the cognitive level in (6), we clarify what is at stake by considering another phenomenon that requires multiple levels of description, namely cipher-based communication. We choose this particular example, because it shares a number of properties with communication through language.

8. The skytale cipher

A cipher relates a readable text (plaintext) to an encrypted text (ciphertext), and can minimally be characterized at two levels, namely as a mathematical function from plaintexts to ciphertexts and as a procedure that performs this mapping. The example we consider here is a system of encryption developed by the Spartans in the fifth century BC. The cryptographic device used is a skytale, a cylinder around which a strip of parchment is wound. Encryption takes place by writing the plaintext message along the length of the cylinder and unwinding the parchment strip. Decryption requires the strip to be wound around a cylinder of identical diameter, after which the plaintext message will reappear. The figure in (18) illustrates the encoding of a short message

using a skytale that allows four lines of text and a length of parchment strip that allows five characters per line. (For ease of exposition we mark empty slots with a dash.)

(18)

T	H	I	S	I
S	T	H	E	P
L	A	I	N	T
E	X	T	-	-

When the strip is unwound, the text is no longer intelligible (we have rotated the letters to increase legibility):

(19)

	T	S	L	E	H	T	A	X	I	H	I	T	S	E	N	-	I	P	T	-	
--	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	--

More abstractly, we can say that the plaintext in (20a) has been encoded as the ciphertext in (20b).

- (20) a. THISISTHEPLAINTEXT--
 b. TSLEHTAXIHITSEN-IPT-

The skytale cipher is one of transposition: it scrambles the letters that make up the plaintext, but does not change them. The scrambling algorithm divides the plaintext into columns and then places each column at the bottom of the one that precedes it. We can describe this algorithm as a function that takes as its input the position of a character in the plaintext (say x) and that returns the position of that character in the ciphertext (say y). What transposition takes place depends on the length of the parchment strip and the diameter of the skytale, as these determine the number of lines (L_{\max}) and the number of characters per line (C_{\max}). The type of transpositions that can be implemented by the procedure sketched above are given by the formula in (21), where L stands for a line of text in the encoding procedure (for (18) the value of L varies from 1 to 4).

(21) $(L - 1) \cdot C_{\max} + 1 \leq x \leq L \cdot C_{\max} \mid y = L_{\max} \cdot (x - ((L - 1) \cdot C_{\max} + 1)) + L$

In order to understand this formula, consider how the position of a character in (18) is related to its position in (19). There are two factors that are relevant. First, the column that contains the character determines the extent to which the character is shifted to obtain its ciphertext position. Second, the character's position in its column remains unchanged. The first factor is expressed by $L_{\max} \cdot (x - ((L - 1) \cdot C_{\max} + 1))$, while the addition of L expresses the constancy of positions internally to 'shifted' columns.

In the case at hand, L_{\max} equals 4 and C_{\max} equals 5. This implies that we can derive from the formula in (21) four more specific formulas, each corresponding to one line. (22a) states that if a character in the plaintext occupies one of the positions 1, 2, 3, 4 or 5, its position in the ciphertext will be as defined by the equation given. (22b–d) give equations for characters whose positions are 6 through 10, 11 through 15, and 16 through 20, respectively.

- (22) a. $1 \leq x \leq 5 \mid y = 4 \cdot (x - 1) + 1$
 b. $6 \leq x \leq 10 \mid y = 4 \cdot (x - 6) + 2$
 c. $11 \leq x \leq 15 \mid y = 4 \cdot (x - 11) + 3$
 d. $16 \leq x \leq 20 \mid y = 4 \cdot (x - 16) + 4$

If the formulas in (22) are applied to positions 1 through 20 in the plaintext, they are mapped to positions in the ciphertext as indicated by the table in (23). This mapping correctly characterizes the transposition in (20).

(23)

1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
1	5	9	13	17	2	6	10	14	18	3	7	11	15	19	4	8	12	16	20

Several aspects of the skytale cipher are relevant to our discussion. To begin with, it is clear that one should distinguish between the mathematical properties of the transposition as given in (21) and (22) and the procedures of encoding and decoding. Suppose that a Martian scientist is presented with a number of plaintexts and their corresponding ciphertexts in the format in (20). He may well discover the cyclical nature of the cipher and come to understand that the key to any cipher of this type consists of two parameters (namely, L_{\max} and C_{\max}). He may subsequently publish a paper in which he proposes the formulas in (21) and (22) and claim that these reflect the reality of the encryption method.

All this would be entirely reasonable, but it is clear from the example at hand that understanding the mathematical properties of the cipher does not identify the actual method of encryption. It merely constrains the class of candidate encoding devices.

It is also clear, when we consider the Spartans who actually used this cipher, that knowing how to use the skytale and the strip of parchment allows one to encode and decode without any knowledge of the formulas in (21) and (22). It is enough that these formulas are ‘built into’ the encryption procedure. The fairest description of the situation is to say that the mathematical properties of the cipher emerge from the encryption and decryption procedures.

The analysis of the skytale cipher exhibits properties typical of information-processing systems. First, the two descriptions to which the cipher yields are equivalent in that both relate the plaintext in (20a) to the ciphertext in (20b). Second, we can say that different algorithms may implement the same cipher. Suppose, for example,

that one encrypts the message in (20a) using a table like (23) rather than a skytale and a strip of parchment. We want to be able to say that this procedure instantiates the same cipher as the one used by the Spartans, but such a generalization can only be made at the level of mathematical description. In other words, the higher level of description abstracts away from certain properties of the lower level. Third, the two levels of description are optimized for their respective functions. While the mathematical description captures the regularities that are the object of study for our Martian scientist, skytale and parchment strip facilitate easy encryption and decryption. Optimization at each level of description results in lack of transparency between these levels. (There is nothing like a mathematical function in the manipulation of stick and strip, or vice versa.)

Consider how the Martian scientist could defend himself against criticism that the mathematical description of the cipher is dispensable, given that it is implied by the methods of encryption and decryption. To begin with, as just mentioned, the mathematical description captures a generalization that cannot be stated at the algorithmic level, because there is more than one algorithm that can compute the same cipher. Furthermore, any communication through a cipher must necessarily yield to a mathematical description of the type in (21) and (22). This is because successful communication relies on the existence of regularities in the association of plaintext and ciphertext and these regularities constitute a description of the cipher at the mathematical level. Finally, the mathematical properties of ciphers are crucial to their selection by code makers. Ciphers are selected on the basis of two factors: they must be difficult to break, but easy to use. The former factor is largely determined by the nature of the mathematical function that describes them, while the latter is a matter of the usability of the associated encryption and decryption devices.

In the case of the skytale cipher, it seems unlikely that its mathematical properties were selected by its users and determined the nature of the encryption procedure. The design of modern ciphers, however, is partly driven by selection for their mathematical properties, which determine, among other things, their breakability. A recent example of selection for higher level properties of a cipher concerns the wish to avoid distribution of keys prior to cipher-based communication. This can be achieved by developing an asymmetric cipher; that is, a cipher which does not use the same key for encryption and decryption. This asymmetry would allow the prospective recipient to make the encryption key public, while not disclosing the decryption key to anyone. Subsequently, the sender can use the public key to encrypt a message that can only be decrypted by the holder of the private key.

Once the idea of an asymmetric cipher was hit upon, the race was on to find a mathematical function that could make it a reality (see Singh 1999 for a vivid description). It will be clear that no one was concerned at this point with the instantiation of this function. In other words, an explanation of the design of the associated encryption and decryption procedures must be partially top-down.

9. Multiple levels of description in language

In the previous section we have clarified why multiple levels of description are required for an information-processing system like a cipher. It is striking that the three properties we have attributed to the skytale cipher also hold of the model in (6).

First, the mapping expressed by the grammar at the cognitive level is logically equivalent to the computations carried out by the performance systems.

Second, the cognitive level is more abstract than the computational and neurological levels. As a result, it generalizes across parser and generator, which are unlikely to employ the same procedures, given their vastly different tasks. This point can be reinforced by a thought experiment. Suppose we build two computers that have radically different underlying architectures that, moreover, employ radically different operating systems. Suppose furthermore that both machines pass a kind of Turing test for language: their performance is indistinguishable from that of native speakers of some natural language. In this situation one would want to say that both computers speak that language, but the notion of identity required for this is unstateable at the hardware level or in terms of the operating system. The relevant notion of identity is based on regularities in the mapping between sound and meaning. The function that specifies these regularities is the grammar.

Third, descriptions at the cognitive and computational level are optimized with respect to their function. The grammar makes explicit regularities in the mapping between sound and meaning in a non-redundant manner, while the performance systems operate with a set of locally applicable rules that may well be highly redundant (such as the sets of rewrite rules referred to in Section 4). This optimization results in a non-transparent relation between grammatical principles and the performance mechanisms from which they emerge. The example used for this throughout the paper is that of Inclusiveness, a principle that captures a fundamental mathematical property of the grammatical code, but which does not lend itself to transparent realization in parser or generator. Therefore, if we restrict the description of language to the computational level, we are unable to express this key principle of human language.

This lack of transparency between the cognitive level and lower levels of description provides a solution to what one might call the paradox of perfection. As Chomsky and others have pointed out, biological systems are typically redundant and robust, while the various components of the grammar have shown to lend themselves to non-redundant analyses (see for example Chomsky & Lasnik 1993:515; and Brody 1998, 2003). Does this pose a problem for the view that grammar is a biological system? We do not think so. The relevant notion of biological system corresponds to a system at a lower level of description. Indeed, as just mentioned, the performance systems are likely to be highly redundant; the same is likely to be true of their physical realization

at the neurological level. But there is no reason why the grammar, which describes a biological system at a higher level of abstraction, should not be perfect or elegant.

This solution to the paradox of perfection shares certain properties with that proposed by Li (1997), who argues that a functional description of a system might be perfect, even if the underlying biological reality is not. Li suggests that our study of grammar is essentially the study of a black box. His claim is that the grammar has the property of perfection because the rules that linguists formulate to relate the input and output of the black box are perfect. On this view, further discoveries about its internal workings may therefore reveal that the proposed rules are an inaccurate or even incorrect description of the underlying biological reality. Clearly, Li's interpretation of grammatical theory is instrumentalist: it is simply a convenient way of talking about the language faculty but does not necessarily capture anything real. (On our reading, Culicover and Nowak (2003) adopt a very similar position.)

There are good, and by now familiar, arguments against an instrumentalist interpretation of grammar. For a start, the thought experiment described above demonstrates that the grammatical description of a language captures a notion that cannot be stated at the algorithmic level. Hence, a growing insight into the biological reality that underlies language cannot replace grammatical theory. (It cannot falsify it either, unless accompanied by an explicit theory about the way in which grammatical properties emerge from the biological level.)

Furthermore, the logic of vocal communication dictates that sound and meaning must be associated in a regular manner and hence yield to a grammatical description. There simply cannot be successful communication in the absence of a shared code. This conclusion holds irrespective of the way this code is computed by speaker and hearer.

Finally, it can be argued that the principles of grammar capture abstract properties of human language that have been shaped by evolutionary pressures. If true, this implies that natural selection is for genes that express themselves in neurological structures that carry out computations that can implement particular (adaptive) grammatical principles (such as Inclusiveness). But if natural selection ultimately targets grammatical principles, then a description of human language at the level of grammar is indispensable. The role of evolution in the design of language is comparable to that of code makers in the development of ciphers. We elaborate on this argument in the next section.

10. Grammar and evolution

What would it take to demonstrate that principles of grammar are shaped by natural selection, and hence not merely instruments that linguists use to talk about the

language faculty? An argument to that effect must begin by establishing that grammar is a complex trait. A simplex trait could be the result of an accidental mutation, but a complex trait must have been shaped by natural selection.

The position that language is indeed a complex trait has been forcefully defended by Pinker and Bloom (1990), Jackendoff (1992, 1994, 2002) and Pinker (1994, 2003). By contrast, Berwick (1998) and Hauser, Chomsky and Fitch (2002) argue it is an extremely simple combinatorial mechanism, inserted as a bridge between pre-existing conceptual abilities and the systems of vocalization. Pinker and Jackendoff (2004) present a detailed critique of Hauser, Chomsky and Fitch's paper, giving empirical arguments that language is special in many more ways than in exhibiting a recursive syntax. We agree with their assessment (see also note 4).

Chomsky (2005) disputes the claim that language is a complex trait on the basis of the archaeological record. About fifty thousand years ago, there seems to have been a sudden cultural spurt, evidenced, among other things, by cave paintings. If this development was triggered by the emergence of human language, there would simply not have been enough time for a complex trait to develop.

The premise on which this argument is based is not itself backed up by evidence. Although cultural developments may be dependent on language, it is not clear that language is a sufficient condition for them. Furthermore, there is archaeological evidence that the evolution of language may have begun much earlier. Martínez et al. (2004) have analyzed the bones of the middle ear belonging to *Homo Heidelbergensis* in order to estimate the range of sound frequencies that they were most sensitive to. A study of five skulls dating back 400,000 years revealed that their owners would have been highly attuned to sounds between three and five kHz. This is remarkably similar to the sensitivity of modern humans, but very different from that of Chimpanzees, our nearest living relatives. It is generally assumed that the sensitivity range in modern humans is due to a specialization for understanding the spoken word. Since *Homo Heidelbergensis* was not one of our direct ancestors, this discovery would trace back the development of language to a shared ancestor who lived about 500,000 years ago. We do not want to suggest that this argument is conclusive, but it does make it clear that the archaeological record is open to more than one interpretation.

The second step in our argument must be to show that principles of grammar contribute to what is adaptive about language. This might seem difficult, as there is disagreement about the primary evolutionary advantage of language. Some researchers claim that it evolved to reap the benefits of communication (see Bickerton 1990; Nowak & Komarova 2001; Pinker 1994, 2003; and Pinker & Bloom 1990), while others support the view that the use of language for communication is secondary and its initial benefit was to allow thinking (conceived of as talking to oneself). Chomsky (2005) is one advocate of this position; he also mentions Jacob (1982) and Tattersall (1998) in this context. However, the two evolutionary scenarios have in common that natural

selection will favour a language faculty that can express more propositional content over one with less expressive power, provided linguistic representations are sufficiently unambiguous to facilitate explicit thought or communication. For example, the ability to encode grammatical dependencies such as θ -role assignment and binding is adaptive in principle, because it makes it possible to achieve interpretational effects that cannot be achieved otherwise. However, if these dependencies are insufficiently constrained, a single structure can be associated with a multitude of interpretations, which would undermine their usefulness. In the extreme case, where any word can be related to any other word in an utterance, the beneficial effects of dependencies are lost altogether. In sum, what is favoured is a language faculty that allows many different grammatical dependencies, as long as these are sufficiently constrained.⁴

As we have argued, the principles that constrain grammatical dependencies (such as Inclusiveness) cannot be transparently realized in the performance systems (for example, as filters). Therefore, the evolutionary pressure described above targets more abstract properties of the language faculty: the grammar described at the cognitive level. If this argument is correct, an instrumentalist interpretation of grammatical theory is untenable.

11. Language acquisition

In the final three sections of this paper we confront our model with some arguments for a grammar module at the same level of description as the performance systems. Perhaps the best-known argument for this position has to do with the complexities of language acquisition. It is often claimed that successful acquisition relies on the setting of a limited number of (usually binary) parameters that together with various universal principles define the space of possible languages.

This view of acquisition is potentially threatened by our view that UG is a description of the language faculty at the cognitive level. If this is true and there is no transparent mapping between cognition and computation, parameters may only be available to the linguist, and not to the language learner. However, in the same way that principles may be non-transparently realized in the parser (and generator), non-transparent realization

4. Note that this line of argumentation also challenges the view that language is a simplex trait. The ‘insertion’ of a combinatory operation will not by itself yield evolutionary benefits, because the ability to express propositional content is adaptive only if it is expressed relatively unambiguously. This requires that any recursive operation co-evolves with a set of constraints, which implies that language is a complex trait.

of parameters is at least a logical possibility. In fact recent work in parameter theory suggests that it is necessary, as we now explain.

Gibson and Wexler (1994) argue convincingly that off-line reasoning about the grammar cannot be part of the process of language acquisition. For this reason, these authors propose a learning strategy according to which a child who fails to parse the current input randomly changes one parameter value and subsequently attempts a reparse. If successful, the new parameter setting is adopted. Otherwise, the initial settings are retained.

In various publications, Janet Dean Fodor and William Sakas (Fodor 1998a,b; Sakas 2000; and Sakas & Fodor 2001) have argued that a significant improvement over this strategy is possible if its randomness can be eliminated. This is possible if off-line inferences about the consequences of parameter setting are built into the acquisition model. More specifically, these authors suggest that universal grammar makes available a set of treelets (or partial trees) that are used in parsing. Parameter setting is conceived of as the selection from this universal set of those treelets that are appropriate to the target language. The child's learning strategy upon encountering an unparseable input, given its current grammar (a language-specific set of treelets), is to select any additional treelets from the universal set that enable it to complete the current parse. These additional treelets are added to the language-specific set if the input is parametrically unambiguous (that is, there are no alternative treelets in the universal set that allow a successful parse). This model of language learning has been shown to be superior to more traditional theories.

Fodor and Sakas's proposal fits in very well with the theory of the competence-performance distinction defended here: saying that universal grammar takes the form of a set of treelets in acquisition and parsing is tantamount to admitting non-transparent realization of grammatical principles. This is because the treelets will not only encode parameters, but must each also adhere to the well-formedness constraints imposed by universal grammar. (In fact, if all treelets are taken to consist of a mother node and two daughters, the proposal is a notational variant of non-transparent realization in the form of rewrite rules.)

We believe a stronger conclusion is warranted: Fodor and Sakas's proposal will require some structuring of the universal set of treelets. The point of parameter theory is to structure the learner's search space in such a way that it is easier to converge on the correct grammar. There are two ways in which parameter setting does so. To begin with, setting a parameter not only implies that a pattern is added to the developing language, but also that alternative patterns are excluded. Furthermore, it might be that the acceptability of more than one pattern depends on a single parameter, so that adding one pattern to the developing language implies adding several other patterns as well (this is the idea behind so-called macro-parameters).

These effects must of course be captured. Even for parameters whose structural effects are quite local and which hence can be successfully expressed as the presence or absence of a particular treelet, it must be assumed that selection of one treelet precludes selection of an alternative. For example, if a treelet is chosen in which the base position of the object follows the verb, then the alternative treelet in which the object precedes it is no longer selectable. If this restriction is not enforced, we cannot truthfully say that we have implemented the OV/VO-parameter.

Even more structure in the set of treelets must be postulated if parameters have non-local effects. Suppose that the OV/VO-parameter not only governs the position of the object with respect to the verb but also the position of other material (such as particles, certain adverbs, resultative secondary predicates, etc.). In that case, the selection of a treelet for one of these constructions must entail automatic selection of other treelets.

We conclude that Fodor and Sakas's proposal, if correct, will require a nontransparent realization of parameter theory in much the same way that parsing requires a nontransparent realization of the grammar.⁵ Just as the abstract nature of linguistic principles interferes with parsing efficiency, the abstract nature of parameters leads to an increase in the computational burden associated with language acquisition. In both cases, the solution seems to lie in an optimization of the language faculty for the computational task at hand and restricting the scope of universal grammar to a more abstract level of description.

12. Aphasia and the shared languages of perception and production

Two further arguments for a grammar module at the same level of description as the performance systems can be found in the following continuation of the quote on p. 2 from Chomsky (2000):

The language faculty has an input receptive system and an output production system, but more than that; no one speaks only Japanese and understands only Swahili. These performance systems access a common body of information, which links them and provides them with instructions of some kind. The performance systems can be selectively impaired, while the cognitive system remains intact [...].

5. One might think that it is not necessary to structure the set of treelets in this way, because the regularities in question are already present in the data that the child is exposed to. It is true that a computer program selecting treelets on the basis of a natural language database would converge on the correct set of treelets, even if the set made available by universal grammar was unstructured. However, if one asks why language displays the regularities that it does, it would beg the question to attribute them to the nature of the data.

We believe that these conclusions are open to criticism.

To begin with, we doubt that there is a pattern of aphasia that conclusively shows that the grammar must be a knowledge module consulted by parser and generator. There are three main cases to consider. First, a person might have lost the ability to produce language, having retained normal perception. On Chomsky's view this would imply that the output production system has been damaged, while the grammar and the input receptive system are unaffected. Second, a person might have lost the ability to perceive language, having retained normal production, implying damage to the input receptive systems, with the grammar and output production system functioning normally. Finally, a person might have lost both language production and perception. This could be indicative of a failure of the grammar, damage to the two performance systems, or both. But these situations can be described equally well if the grammar is not a module at the computational level: the first pattern of aphasia corresponds to a breakdown of the generator, the second to a breakdown of the parser, and the third to a breakdown of both.

The second point in the quote above concerns the question why the languages of perception and production are systematically identical. Chomsky's answer is that production and comprehension make use of the same knowledge base. But this does not solve the puzzle. After all, if it were adaptive for the languages of comprehension and production to be different, evolution could very well have led to separate grammars for the two systems.

In fact, there are obvious arguments suggesting evolutionary pressure towards a common language for perception and production, whether language is conceived as a tool for communication or for thought. This is compatible with the view that parser and generator share a knowledge base, but it does not force us to adopt this position. For the sake of argument, let us assume that the grammar employed in perception is constant across individuals. The perception system is a crucial component of the evolutionary environment to which the production system is sensitive. Hence, the production system must have evolved to produce utterances that optimally fit the perception system. The same line of argumentation implies that the perception system must have evolved in such a way that it optimally processes the output of the production system. Therefore, evolutionary negotiation will result in parser and generator settling on the same code.

Of course, there is variation between languages and this means that the above argument merely guarantees that parser and generator adhere to the same universal grammar. However, the same evolutionary pressures favour the development of mechanisms that ensure identical parameter setting in parser and generator. There is a very likely candidate for such a mechanism, namely feedback in production: hearing oneself speak. There is evidence that monitoring one's speech is extremely important in production. A subject wearing headphones that delay this auditory feedback by

200 msec. or so displays severely disrupted speech. He or she will sometimes even produce uncontrollable stuttering (see Jackendoff 1987 and references mentioned there). In light of this, it is suggestive that language perception systematically precedes language production in children.

13. Grammaticality judgments

A final argument for the existence of a competence module concerns the distinction that native speakers make between sentences that are hard to process and those that are ungrammatical. That speakers can make this distinction is all the more remarkable in view of the fact that giving a grammaticality judgment is a matter of performance. In order to know that a mistake has been made, there has to be some yardstick against which performance is judged and the usual claim is that this is the role that the competence module plays.

Given the proposal we have put forward, the question arises how we could characterize the relevant contrast. Of course, we accept that there is a difference between ungrammatical sentences and sentences that are hard to parse. We also accept that native speakers have a certain degree of access to this distinction. What we do not accept is that this distinction must be grounded in a competence module at the computational level. Recall that we have taken the parser and generator to contain a non-transparent realization of the grammar. Hence, the performance systems will not be able to assign a complete structure to strings that do not adhere to the hearer's grammatical code. This situation must be distinguished from one in which it is hard – but not impossible – to find a complete structure for the input string, such as in the case of garden path sentences. Native speakers can make the distinction between these two cases by producing the string in question several times and using it as input to the parser through the feedback mechanism discussed above. This 'looping' may reveal that a structure is in fact available although it was not found on the first attempt.

As is well-known, a native speaker may still assign an interpretation to an input for which no complete structure is found. In particular, if the input string can be parsed into several unconnected structures, discourse-level processes may be able to integrate the semantics of these fragments into a coherent proposition. (It is probably not an overstatement that this type of processing is very common in normal conversations.) For certain ill-formed sentences it is also conceivable that the performance mechanisms carry out repairs that reconcile the input with the grammatical code. A low-level example of such a repair is phoneme restoration (Warren 1970), but there might be higher-level repairs that locally affect phoneme, morpheme or constituent order. In this respect, there is no difference between our proposal and the standard view.

One aspect in which our proposal does diverge from the standard view is that it has something to say about the well-formedness of incomplete structures. The competence grammar, conceived of as a filtering device, will rule out all incomplete structures. However, such structures are frequently produced in discourse and not perceived as deviant, beyond the fact that they are incomplete. Someone listening to the monologue in (24), for example, will not experience the various incomplete utterances as ill-formed.

- (24) I said... I mean... Susan's unlikely to do something like that. What d'you think she... What would you say if she'd really done that?

By contrast, irrespective of the context in which it is uttered, (25) is perceived as not just incomplete but deviant (see Jackendoff 1987: 106).

- (25) What movies do you know lots of people who...

Our view of the performance mechanisms is that they are designed to operate in accordance with a grammatical code. However, as opposed to the competence grammar, such mechanisms must be able to deal with incomplete structures, as intermediate stages in the processes of parsing and generation are by necessity incomplete.⁶ To put the same thing differently, in normal circumstances the performance mechanisms operate in accordance with the grammatical code and hence any partial structures they assign are reflections of that code. This implies that the judgement associated with (25) signals that the perceptual mechanisms cannot find a single structure that covers the entire string. Since no such failure occurs in any of the incomplete structures in (24), these are not considered deviant.

Standard grammaticality judgement tasks always involve judgements of strings that are taken to form complete sentences. Therefore, the task comes with the instruction to check whether a single complete structure can be assigned to the input. While such a task can reveal a great deal about the language faculty, the fact that it targets an idealized language means that what is acceptable in conversation will fail in the grammaticality judgment task: speakers will judge both the incomplete WH questions in (24) and (25) ungrammatical.

The data in (24) and (25) are not completely beyond the reach of theories that assume a grammatical module (conceived of as a filtering device), provided that a further module is added to the language faculty whose task it is to uncover possible continuations of incomplete structures. The resulting completed structures can be presented to the grammar in normal conversation, but the sentence completion module

6. See Ristad (1993: 115–116) for discussion.

would have to be switched off in grammaticality judgment tasks. This will work, but it seems ad hoc.

References

- Abney, S. 1988. On the notions 'GB Parser' and 'Psychological Reality'. In *MIT Parsing Volume*, 1987–88, S. Abney (ed). Cambridge MA: Center for Cognitive Science, MIT.
- Abney, S. 1989. A computational model of human parsing. *Journal of Psycholinguistic Research* 18: 129–144.
- Abney, S. 1991. Parsing by chunks. In *Principle-Based Parsing: Computation and Psycholinguistics*, R. Berwick, S. Abney & C. Tenny (eds), 257–278. Dordrecht: Kluwer.
- Berwick, R. 1998. Language evolution and the minimalist program: The origins of syntax. In *Approaches to the Evolution of Language: Social and Cognitive Bases*, J. Hurford et al. (eds), 320–340. Cambridge: CUP.
- Berwick, R., Abney, S. & Tenny, C. (eds). 1991. *Principle-Based Parsing: Computation and Psycholinguistics*. Dordrecht: Kluwer.
- Berwick, R. & Weinberg, A. 1984. *The Grammatical Basis of Linguistic Performance*. Cambridge: MIT Press.
- Berwick, R. & Weinberg, A. 1985. Deterministic parsing and linguistic explanation. *Language and Cognitive Processes* 1: 109–134.
- Bickerton, D. 1990. *Language and Species*. Chicago IL: University of Chicago Press.
- Brody, M. 1998. The minimalist program and a perfect syntax. *Mind and Language* 13: 205–221.
- Brody, M. 2003. *Towards an Elegant Syntax*. London: Routledge.
- Chomsky, N. 1965. *Aspects of the Theory of Syntax*. Cambridge MA: The MIT Press.
- Chomsky, N. 1995a. Bare phrase structure. In *Government and Binding Theory and the Minimalist Program*, G. Webelhuth (ed.), 383–439. Oxford: Blackwell.
- Chomsky, N. 1995b. *The Minimalist Program*. Cambridge MA: The MIT Press.
- Chomsky, N. 2000. *New Horizons in the Study of Language and Mind*. Cambridge: CUP.
- Chomsky, N. 2005. Three factors in language design. *Linguistic Inquiry* 36: 1–22.
- Chomsky, N. & Lasnik, H. 1993. The theory of principles and parameters. In *Syntax: An International Handbook of Contemporary Research*, J. Jacobs et al. (eds), 506–569. Berlin: Walter de Gruyter.
- Culicover, P. & Nowak, A. 2003. *Dynamical Grammar*. Oxford: OUP.
- Dennett, D. 1992. *Consciousness Explained*. London: Penguin books.
- Earley, J. 1970. An efficient context-free parsing algorithm. *Communications of the Association for Computing Machinery* 14: 453–460.
- Fodor, J.D. 1998a. Unambiguous triggers. *Linguistic Inquiry* 29: 1–36.
- Fodor, J.D. 1998b. Parsing to learn. *Journal of Psycholinguistic Research* 27: 339–374.
- Fong, S. 1991. The computational implementation of principle-based parsers. In *Principle-Based Parsing: Computation and Psycholinguistics*, R. Berwick, S. Abney & C. Tenny (eds), 65–82. Dordrecht: Kluwer.
- Frazier, L. & Clifton, C. 1995. *Construal*. Cambridge MA: The MIT Press.
- Gibson, E. & Wexler, K. 1994. Triggers. *Linguistic Inquiry* 25: 407–454.
- Gorrell, P. 1995. *Syntax and Parsing*. Cambridge: CUP.
- Hauser, M., Chomsky, N. & Fitch, T. 2002. The faculty of language: What is it, who has it, and how did it evolve? *Science* 298: 1569–1579.

- Jackendoff, R. 1987. *Consciousness and the Computational Mind*. Cambridge MA: The MIT Press.
- Jackendoff, R. 1992. *Languages of the Mind*. Cambridge MA: The MIT Press.
- Jackendoff, R. 1994. *Patterns in the Mind: Language and Human Nature*. New York NY: Basic Books.
- Jackendoff, R. 2002. *Foundations of Language: Brain, Meaning, Grammar and Evolution*. Oxford: OUP.
- Jacob, F. 1982. *The Possible and the Actual*. Washington DC: University of Washington Press.
- Kempson, R., Meyer-Viol, W. & Gabbay, D. 2001. *Dynamic Syntax: The Flow of Language Understanding*. Oxford: Blackwell.
- Kolb, H.-P. & Thiersch, C. 1991. Levels and empty categories in a principles and parameters approach to parsing. In *Representation and Derivation in the Theory of Grammar*, H. Haider & K. Netter (eds), 251–302. Dordrecht: Kluwer.
- Koot, H. van de 1990. *An Essay on Grammar-Parser Relations*. Dordrecht: Foris.
- Levelt, W. 1989. *Speaking: From Intention to Articulation*. Cambridge MA: The MIT Press.
- Li, Y. 1997. An optimized universal grammar and biological redundancies. *Linguistic Inquiry* 28: 170–178.
- Marcus, M. 1980. *A Theory of Syntactic Recognition for Natural Language*. Cambridge MA: The MIT Press.
- Marcus, M. 1987. Deterministic parsing and description theory. In *Linguistic Theory and Computer Applications*, P. Whitelock, H. Somers, P. Bennett, R. Johnson & M. McGee Wood (eds), 69–112. New York NY: Academic Press.
- Marcus, M., Hindle, D. & Fleck, M. 1983. D-Theory: Talking about talking about trees. *Association for Computational Linguistics* 21: 129–136.
- Marr, D. 1982. *Vision*. New York NY: W.H. Freeman.
- Martínez, I. et al. 2004. Auditory capacities in middle pleistocene humans from the Sierra de Atapuerca in Spain. *Proceedings of the National Academy of Sciences of the United States of America* 101: 9976–9981.
- Neeleman, A. & van de Koot, H. 2002. The configurational matrix. *Linguistic Inquiry* 33: 529–574.
- Newmeyer, F. 2003. Grammar is grammar and usage is usage. *Language* 79: 682–707.
- Nowak, M.A. & Komarova, N.L. 2001. Towards an evolutionary theory of language. *Trends in Cognitive Sciences* 5(7): 288–295.
- Phillips, C. 1996. Order and Structure. Ph.D. dissertation, MIT.
- Phillips, C. 2003. Order and constituency. *Linguistic Inquiry* 34: 37–90.
- Pinker, S. 1994. *The Language Instinct*. New York NY: Harper Collins.
- Pinker, S. 2003. Language as an adaptation to the cognitive niche. In *Language Evolution: States of the Art*, M. Christiansen & S. Kirby (eds), 16–37. Oxford: OUP.
- Pinker, S. & Bloom, P. 1990. Natural language and natural selection. *Behavioral and Brain Sciences* 13: 707–784.
- Pinker, S. & Jackendoff, R. 2004. The faculty of language: What's special about it? Ms, Harvard University & Brandeis University. To appear in *Cognition*.
- Popper, K. 1992. *Quantum Theory and the Schism in Physics: From the Postscript to the Logic of Scientific Discovery*. London: Routledge.
- Pritchett, B. 1992. *Grammatical Competence and Parsing Performance*. Chicago IL: The University of Chicago Press.
- Ristad, E. 1993. *The Language Complexity Game*. Cambridge MA: The MIT Press.
- Sakas, W.G. 2000. Modeling the effect of cross-language ambiguity on human syntax acquisition. In *Proceedings of CoNLL-2000 and LLL-2000*, 61–66. Lisbon.

- Sakas, W.G. & Fodor, J.D. 2001. The structural triggers learner. In *Language Acquisition and Learnability*, S. Bertolo (ed.), 172–233. Cambridge: CUP.
- Singh, S. 1999. *The Code Book: The Secret History of Codes and Code-breaking*. London: Fourth Estate.
- Steedman, M. 2000. *The Syntactic Process*. Cambridge MA: The MIT Press.
- Tattersall, I. 1998. *The Origin of the Human Capacity*. New York NY: American Museum of Natural History.
- Warren, R. 1970. Perceptual restoration of missing speech sounds. *Science* 167: 392–393.
- Younger, D. 1967. Recognition and parsing of context-free Languages in time n^3 . *Information and Control* 10: 189–208.

Monitoring for speech errors has different functions in inner and overt speech

Sieb Nooteboom
Utrecht University

In this paper it is argued that monitoring for speech errors is not the same in inner speech and in overt speech. In inner speech it is meant to prevent the errors from becoming public, in overt speech to repair the damage caused by the errors. It is expected that in inner speech, but not in overt speech, more nonword errors are detected than real-word ones, and that overt repairs of errors detected in inner speech differ from overt repairs of errors detected in overt speech in that they have shorter offset-to-repair times, are spoken with raised instead of lowered intensity and pitch, and are less often accompanied by editing expressions. These hypotheses are tested against a collection of experimentally elicited spoonerisms and a collection of speech errors in spontaneous Dutch. The hypotheses are basically confirmed.

1. Introduction

Baars, Motley and MacKay (1975) suggested that in speech production there is pre-articulatory editing of inner speech, during which speech errors are detected and repaired before they surface in overt speech. This assumption was needed to explain the phenomenon of so-called lexical bias in segmental speech errors. Lexical bias is the phenomenon that segmental speech errors more often create real words than non-words, other things being equal. Lexical bias has been attested both in experimentally elicited speech errors (e.g. Baars et al. 1975; Dell 1986; Nooteboom 2005b; Nooteboom & Quené 2008) and in spontaneous speech errors (e.g. Dell & Reich 1981; Nooteboom 2005a; Hartsuiker, Antón-Méndez, Roelstraete & Costa 2006; but see Garrett 1976 and Del Viso, Igoa & Garcia-Albea 1991). The idea is that nonword errors are more often detected, rejected and repaired in inner speech than real-word errors, causing a greater frequency of real-word errors in overt speech. The need for such covert, rapid and fluent pre-articulatory editing was rejected by Stemmer (1985) and Dell (1986). They explained lexical bias from immediate feedback of activation between phonemes and lexical units within the mental speech production system proper. However, Levelt (1989) and Levelt, Roelofs and Meyer (1999) rejected feedback, and revived the original explanation by Baars et al. More recently, Hartsuiker, Corley and Martensen (2005) provided

evidence that lexical bias in segmental speech errors is caused both by immediate feedback of activation between speech sounds and word forms, as proposed by Dell (1986), and by covert monitoring of inner speech, applying a criterion of lexicality, as proposed by Levelt et al. (1999). This was further confirmed by Nooteboom and Quené (2008) who in some conditions demonstrated a positive lexical bias in early interrupted speech errors such as *b.. dark boat*, but in other conditions found a reverse lexical bias in such early interruptions. Importantly, interruptions as these must have been reactions to inner speech, because speech fragments like *b..* are shorter than a humanly possible reaction time. Therefore the speaker cannot have reacted to detecting the error in his overt speech. The positive lexical bias in interrupted speech errors cannot have been caused by covert monitoring, because repairs of such interrupted speech errors are made overtly, not covertly, and thus such positive lexical bias provides evidence for feedback. The reverse lexical bias that was also found cannot be explained from feedback, because feedback of the kind proposed by Stemberger (1985) and Dell (1986) never leads to a reverse lexical bias. Such reverse lexical bias thus provides evidence for an active monitor, detecting and rejecting nonwords in inner speech more frequently than real words. These findings bring into focus the psychological reality of an active monitor, covertly detecting, rejecting and sometimes repairing speech errors in inner speech. Apparently, lexical bias in segmental speech errors results both from feedback of activation between speech sounds and word forms in the mental production of speech, and from a criterion of lexicality used in monitoring inner speech for speech errors. It should be said, though, that this lexicality criterion only modulates a relatively high level of detection of both real-word errors and nonword errors. Obviously, error detection does not depend exclusively on a lexicality criterion (cf. Nooteboom & Quené 2008).

Kolk (1995) has suggested that if at the moment of repair the word's lemma is still active, it is sufficient to start the compilation of the form once more. This would explain the rapidity of the repair. The implied alternative would be that the correct target is accessed for a second time all the way from the monitor which is supposed to reside at the conceptual level (cf. Levelt et al. 1999). But in Kolk's formulation it is obviously assumed that the correct target was not available at the moment of error detection. If so, it remains unclear how segmental real-word errors can be rapidly detected. It seems more reasonable to assume that error form and correct target are simultaneously available and competing for the same slot in inner speech. Normally, the most activated form would win, but then the monitor can detect the difference and decide that an error has been made. The assumption that error form and correct target are simultaneously competing for the same slot in inner speech is supported by the relatively frequent occurrence of speech errors such as *gfgfg..feit goud* in experiments eliciting consonant exchanges (Nooteboom 2007). Obviously in this case *feit* and *goud* are competing for the same slot. Also I found a few interesting cases where the participant very softly

whispered the correct target and then immediately, with no silent interval, spoke the elicited spoonerism aloud. This also suggests that correct target and speech error are competing for the same slot in inner speech.

At first sight the supposed monitoring of inner speech for speech errors looks very similar to what happens in monitoring overt speech for speech errors. Such overt errors are not always but in c. 50% of the cases detected and then overtly repaired by the speaker (Nooteboom 1980). This leads to such sequences as *bark boat.. dark boat*. Indeed Levelt (1989) and Levelt et al. (1999) have suggested that there is only one monitor checking speech for errors. This monitor employs the same speech comprehension system that is also involved in perceiving other-produced speech. The comprehension system receives two forms of input, one being inner speech before it is articulated. This closes the internal perceptual loop. The other input is overt speech processed by the auditory system. This closes the external perceptual loop. The suggestion is that, apart from the input, the monitoring process is identical in both cases.

The idea that monitoring of inner and monitoring of overt speech are identical also seems implied by three prerequisites for a satisfactory account of any monitoring bias, formulated by Hartsuiker (2006), who stated as prerequisites that (1) the proposed account poses functional monitoring criteria; (2) the bias can be altered by manipulations affecting monitoring performance; (3) the monitoring bias occurs also in perception. The third prerequisite suggests that the same bias that is assumed to operate in monitoring inner speech should also be found in speech perception, and of course, one can perceive both speech produced by others and speech produced by oneself. If one thinks of the latter, Hartsuiker's third prerequisite suggests that if one assumes that monitoring inner speech employs a criterion of lexicality causing lexical bias in overt speech errors, then one would predict that in monitoring overt speech for speech errors also more nonwords than real words are detected, rejected and repaired. Hartsuiker also predicts, of course, that a similar criterion operates in the perception of other-produced speech. The suggestion clearly is that finding a lexical bias in the detection of overt speech errors, either produced by the speaker himself or by another speaker, would support the assumption of a lexicality criterion being employed by the monitor operating on inner speech.

However, it might be argued that self-monitoring of inner speech and self-monitoring of overt speech have different functions and different time-constraints. If so, it might be expected that the strategies involved are not necessarily identical, and this may lead to different behavioural patterns. Monitoring inner speech for speech errors probably attempts to prevent errors in inner speech from becoming public, preferably without any perceivable effects on the fluency of the speech produced, or at least with so little perceivable effect as possible. For this reason, it is reasonable to assume that detecting and repairing errors in inner speech is under some time pressure. Monitoring overt speech for speech errors obviously comes too late to prevent those errors from becoming public. There is no reason to suppose that detecting and repairing overt speech errors

is under time pressure. Rather the speaker should take his or her time to make clear to the listener that an error has been made and to prevent damage to communication by this error. These considerations lead to some testable predictions.

If indeed monitoring for speech errors in inner speech strives towards making the error as little noticeable as possible for the listeners, one would expect that speakers hasten to repair the error. A speaker may repair the detected error before it is spoken at all. If so, there is no overt repair. But, probably due to the time constraints on monitoring, rather often a speaker inadvertently initiates speaking the error detected in inner speech, and then rapidly interrupts his or her speech. This leads to errors of the type *s..fat soap*. The basic assumption here is that both the interruption and the following repair are intended to make the error as little noticeable as possible. This leads to time pressure. This is potentially different for errors that are only detected after they have been spoken. The full error has been made and very likely also heard by the listener. In this case it is more relevant for the speaker to let the listener know that an error has been made and that in the interpretation the error has to be replaced by the repair. Here there is not necessarily any time pressure on the speaker. Now it has recently been shown that under time pressure, other things being equal, (unless error and target are phonetically very similar) nonword errors are more often detected than real-word ones, but if the time pressure is removed nonword errors are equally often detected as real-word ones (Nooteboom & Quené 2008). If this result is valid, one expects more nonword errors to be detected than real-word ones in normal monitoring of inner speech, but not in normal monitoring of overt speech. So here is a first hypothesis to be tested.

A second hypothesis is concerned with the offset-to-repair times. If indeed speakers hasten to cover up any errors that are detected in inner speech but yet inadvertently initiated, as in *s..fat soap*, one expects offset-to-repair times to be generally short in such errors, whereas offset-to-repair times following errors detected in overt speech are probably much longer, because in these cases the speaker has no reason to hurry, and may need time to plan the repair in such a way that the listener knows that an error has been made and should be replaced with the repair in order to arrive at the intended interpretation. Of course, as has been pointed out by Blackmer and Mitton (1991), if the error is detected in inner speech and yet, inadvertently, spoken as an initial fragment, not only the interruption but also the repair may have been planned before speech was initiated. In such cases one may find offset-to-repair times of 0 ms.

A third hypothesis is concerned with speech prosody. If indeed speakers have a tendency to cover up the overt consequences of errors detected in inner speech, they may be expected not only to have very short offset-to-repair times, but also to make the inadvertent error fragment less noticeable by speaking the rapidly following repair with more intensity and on a higher pitch than the error fragment itself was spoken. For errors detected in overt speech the situation is different. It is too late for any cover up anyway, and the error may potentially lead to a wrong interpretation by the listener.

Thus here the error must remain relatively noticeable, and the repair should stand out as a repair by a prosody that is markedly different from both the error and the regular correct responses. This could be achieved by speaking the repair with somewhat less intensity and on a lower pitch than the error itself.

Finally, a fourth hypothesis relates to the editing expressions, such as *sorry*, *oh!*, *er*, *correction*, or *hahaha*, that may accompany a repair. If speakers tend to cover up the overt consequences of errors detected in inner speech, they very likely will not use such editing expressions in their overt repairs of such errors. But in making repairs of errors detected in overt speech, such editing expressions would be functional (Levelt 1983), and one may thus expect the frequent use of such editing expressions.

From the assumption that monitoring inner speech and monitoring overt speech have different functions, one can thus derive several hypotheses relating to different behavioural patterns for the two types of repair. However, to test these hypotheses, it should be clear which overt repairs are reactions to error detection in inner speech and which are reactions to error detection in overt speech. It is not self-evident that this is always easy to do. Although in cases like *s..fat soap* it is very unlikely that the interruption results from error detection in overt speech in view of the very short duration of the erroneous speech fragment, in cases with interruptions that come somewhat or much later it is not clear whether these are relatively late reactions to errors in inner speech or whether these are reactions to errors in overt speech, or perhaps to both simultaneously. Hartsuiker and Kolk (2001), using a computational model of self-monitoring including both an internal and an external perceptual loop, have demonstrated that empirical distributions of overt offset-to-repair times cannot be simulated satisfactorily with only an external loop. A considerable contribution of the internal loop is needed in order to get a reasonable fit of the empirical distributions of overt offset-to-repair times. Hartsuiker, Kolk and Martensen (2005), using an equation with two unknowns for estimating the relative numbers of detected and undetected speech errors in inner from those in overt speech, found that the accuracy of the internal perceptual loop is considerably better than that of the external perceptual loop. The accuracy of the latter may in some conditions even be zero. The corollary of this is that many overt interruptions and repairs are reactions to the detection of errors in inner speech. However, it is reasonable to assume that the later the interruption comes after the error has been made the greater the probability that the error was detected in overt speech. Below we will assume that early interrupted speech errors of the type *s..fat soap* or *sa..fat soap* are errors detected in inner speech and that repaired completed speech errors of the type *sat soap... fat soap* or *sat foap...fat soap* are errors detected in overt speech (for evidence that this assumption at least is highly plausible see Nooteboom 2005b).

Obviously, that an interruption was planned in inner speech does not necessarily mean that the repair was also planned in inner speech. However, if one finds an early interruption such as *s..fat soap* with an offset-to-repair time of zero ms, or not much

longer, it seems almost certain that the repair was planned before speech was initiated (Blackmer & Mitton 1991). Experimental evidence that often repairs are planned before the speech is interrupted is obtained by Hartsuiker, Catchpole, De Jong and Pickering (2008). But if the offset-to-repair time is many hundreds of ms, it is not unlikely that, although the interruption was already planned before speech initiation, the repair was only planned after the overt speech fragment was spoken (and perhaps re-detected in overt speech). Likewise, if there is a repaired completed exchange with an offset-to-repair time of 0 ms, this may be a case where the error had been detected and the repair had already been planned before speech initiation, whereas with much longer offset-to-repair times this is much less likely. The point here is that if repairs of early interruptions are interpreted as reactions to error detection in inner speech, and repairs to completed speech errors as reactions to error detection in overt speech, this should be seen as probabilistic. It seems unlikely that a 100% separation between repairs as reactions to error detection in inner speech and repairs as reactions to error detection in overt speech can be achieved. Some repairs of interrupted speech errors may be reactions to overt speech (although the interruptions themselves were reactions to error detection in inner speech) and some repairs of completed speech errors may be reactions to error detection in inner speech. This implies that individual cases have little to say about the validity of the earlier hypotheses. These can only be tested statistically.

Before the earlier mentioned hypotheses are put to the test, it may be good to go back for a moment to the original hypothesis by Baars et al. (1975), viz. that there is pre-articulatory editing of inner speech, covertly and fluently detecting, rejecting and repairing speech errors before they are spoken. Obviously such cases are not observable. This makes it difficult to ascertain the reality of covert and fluent editing of inner speech. Yet there is at least one convincing demonstration of such covert and fluent editing. This is found in data published by Motley, Camden and Baars (1982) who demonstrated activity of the monitor by measuring Galvanic Skin Responses to experimentally elicited but suppressed taboo words in otherwise perfectly fluent and correct speech. The earlier mentioned model by Hartsuiker and Kolk (2001) allows such covert and fluent editing. However, Hartsuiker, Kolk and Martensen (2005) assume that, although sometimes covert repairs may leave no observable traces, often covert repairs lead to observable disfluencies. Unfortunately it is not easy to investigate unobservable covert repairs, and it is also not easy to know whether specific disfluencies stem from covert repairs of speech errors or from other causes. Therefore the remainder of this paper will focus on overt interruptions and repairs, being either reactions to speech errors in inner speech or to speech errors in overt speech. This will be done both for experimentally elicited speech errors and for spontaneous speech errors.

The following four hypotheses will be tested, under the assumption that most early interruptions and their repairs are reactions to error detection in inner

speech and most repaired completed speech errors are reactions to error detection in overt speech:

1. Other things being equal, more potential nonword than potential real-word errors are early interrupted (Nootboom 2005b; Nootboom & Quené 2008). However, there are equally many nonword and real-word repaired completed speech errors.
2. Overt repairs following interrupted speech errors have shorter offset-to-repair times than overt repairs following completed speech errors.
3. Overt repairs following interrupted speech errors have more intensity and higher pitch than the errors themselves. Overt repairs following completed speech errors have less intensity and lower pitch than the errors themselves.
4. Overt repairs following interrupted speech errors are much less often accompanied by editing expressions than overt repairs following completed speech errors.

2. The data: Elicited speech errors and errors in spontaneous speech

The speech errors and their interruptions and repairs used in this paper have been described in earlier publications, but the relevant patterns in the data described below are mostly new. The corpus of speech errors in spontaneous Dutch was described in Nootboom (2005a). The corpus of experimentally elicited Dutch speech errors was described partly in Nootboom (2005b) and partly in Nootboom and Quené (2008).

2.1 Elicited speech errors

These errors stem from classical SLIP (Spoonerisms of Laboratory-Induced Predisposition) experiments, meant to elicit consonant exchanges like *fine book* becoming *bine fook* (nonword error) or *cool tap* becoming *tool cap* (real-word error). The SLIP technique basically works as follows: Subjects are successively presented visually, for example on a computer screen, with word pairs such as *dove ball*, *deer back*, *dark bone*, *barn door*, to be read silently. On a prompt, for example a buzz sound or a series of question marks (“?????”), the last word pair seen (the test word pair as opposed to the biasing word pairs), in this example *barn door*, has to be spoken aloud. Interstimulus intervals are in the order of 1000 ms, as is the interval between the test word pair and the prompt to speak. Every now and then a word pair like *barn door* will be mispronounced as *darn bore*, as a result of segmental biasing by the preceding word pairs.

Below the interrupted elicited exchanges such as *da..barn door*, are interpreted as reactions to error detection in inner speech, the repaired completed elicited exchanges such as *barn door* >> *darn bore* are interpreted as reactions to error detection in overt

speech. The particular selection of cases used to test the hypotheses is explained and argued separately for each hypothesis.

2.2 Speech errors in spontaneous speech

The collection of speech errors in spontaneous speech described in Nootboom (2005a) in fact consists of two separate corpora. One of these, the AC/SN corpus, does not contain any repairs, and therefore is useless for the current purposes. The other, the Utrecht corpus, contains speech errors together with any repairs, noted down between 1977 and 1982 by staff members of the Utrecht Phonetics department on the initiative of the late Anthony Cohen, at that time professor of Phonetics in Utrecht. Of the 2500 errors in the collection 1100 are segmental syntagmatic errors. These will be used here. As these errors were noted down in standard orthography, or sometimes in phonetic script, and the sound has not been recorded, obviously there are no known offset-to-repair times. The moment of stopping for making a repair can of course be estimated from the amount of speech material spoken after the error has been made. It was found that after a segmental error only very rarely a speaker speaks more than the error form itself before stopping. This is different for lexical errors where in many cases the speaker stops only after a few more words; cf. Nootboom (2005a). So we can distinguish between two classes of repairs: The class of repairs following an early interruption of the error form, as in *b.. thicker bush* or *bi.. thicker bush*, and a class of repairs following the completed error form as in *bicker...thicker bush*. All other cases are so rare that they are statistically irrelevant.

3. Testing the hypotheses

3.1 Hypothesis 1: *Other things being equal, there are more nonword than real-word errors early interrupted. There is no such difference between nonword and real-word repaired completed speech errors.*

This hypothesis will first be tested on elicited speech errors and then on spontaneous speech errors.

3.1.1 Elicited speech errors

For the current purpose the speech errors are taken from two published experiments as reported in Nootboom (2005b) and in Nootboom and Quené (2008). The first publication describes only a single classical SLIP experiment, the latter described two such experiments. However, of these two experiments for testing this hypothesis only the data of Experiment 1 will be used. The reason is that in Experiment 2 time pressure

was on purpose artificially removed, whereby the conditions for normal monitoring of inner speech were changed. In order to obtain a reasonably large set of experimental data the data of Nooteboom (2005b) and Experiment 1 of Nooteboom and Quené (2008) are collapsed. These experiments are in all relevant aspects comparable.

As earlier indicated, each early interruption, whether or not followed by an overt repair, is interpreted as a reaction to error detection in inner speech and each overt repair of a completed consonant exchange is interpreted as a reaction to error detection in overt speech. It should be noted that interruption itself marks an error as detected, whether or not it is repaired later. For the completed exchanges in a SLIP experiment, the only evidence of detection by the monitor is the presence of a repair. For the main pattern in the data see Table 1.

Table 1. Numbers of unrepaired completed exchanges, early interruptions, and repaired completed exchanges, from two SLIP experiments combined (see text)

	unrepaired completed exchanges	interruptions	repaired completed exchanges
Real-word errors	74	65	19
nonword errors	41	104	10

The assumption made here is that unrepaired completed exchanges have not been detected by the speaker at all, interruptions have been detected in inner speech, and repaired completed exchanges have been detected in overt speech. The point of interest is whether in inner speech, in overt speech, or in both, nonword errors are detected more frequently than real-word ones. To find out, the distributions of real-word and nonword errors for interruptions and for repaired completed errors are compared with the same distribution for the undetected errors. Obviously, there may be uncertainty about the lexicality of an error when it has been early interrupted. However, these errors were made in two conditions, one condition eliciting real-word errors, the other eliciting nonword errors. It is assumed here that the interrupted errors conform to the eliciting condition. The data clearly show that the distribution over real-word and nonword errors for the interruptions differs significantly from the same distribution for undetected errors (Fisher's exact test: $p < .0001$). The distribution for interruptions can only be explained by assuming that a lexicality criterion has been operative in monitoring inner speech. The data also show that the distribution over real-word and nonword errors for the repaired completed errors does not differ from the same distribution for undetected errors (Fisher's exact test: $p = 1$). This can only be explained by assuming that in monitoring overt speech no lexicality criterion is applied. As predicted, early interruptions reflect a criterion of lexicality, repairs of completed exchanges do not.

3.1.2 *Speech errors in spontaneous speech*

To find out whether the pattern in experimentally elicited errors is supported by the pattern in speech errors made in spontaneous speech, a count was made in the Utrecht corpus of speech errors of how many segmental speech errors were interrupted before completion of the error form, separately for errors that, when completed, would have generated a real word and for errors that would have generated a nonword. It was also counted how often the error form was completed but the speaker had stopped after the error form for making a repair. In the collection there are 1100 segmental speech errors. After removal of all those cases that for a variety of reasons were difficult to classify, there remained 744 segmental errors.

Table 2. Numbers of unrepaired completed segmental speech errors, early interruptions, and repaired segmental speech errors, taken from a collection of speech errors made in spontaneous Dutch (see text)

	unrepaired completed exchanges	interruptions	repaired completed exchanges
Real-word errors	163	52	121
nonword errors	158	127	123

The relevant classification of these errors, comparable to the classification in Table 1, can be seen in Table 2. Again, as with the elicited speech errors, unrepaired completed exchanges are interpreted as not being detected by the speaker, interruptions are interpreted as detected in inner speech, and repaired completed exchanges as errors detected in overt speech. It might be observed that the numbers of unrepaired completed exchanges do not seem to show much of a lexical bias. However, this is misleading. The probability of a real-word error to be made is roughly 50% for CVC words, but rapidly decreases with word length. In the corpus words of all possible lengths were included. Average length of the words in the corpus that were involved in segmental speech errors is roughly 6 phonemes. The lexical bias must be considerable if the numbers for real-word and nonword errors are about equal. Again, as earlier with the elicited errors, there may be uncertainty about the lexicality of interrupted errors. An English example might be *ba.. marvelous boat*. In each individual interruption the lexicality was judged from the combination of speech error and repair. In this example, the interrupted speech error is assumed to be *barvelous*, which is a nonword. It may be seen that the pattern of the data in Table 2 is very similar to the pattern found for elicited speech errors in Table 1. The distribution over real-word and nonword errors for the interruptions differs significantly from this distribution for the undetected completed speech errors (Fisher's exact test: $p < .0001$). The distribution for the interruptions suggests that in monitoring inner speech a criterion of lexicality is applied.

The distribution over real-word and nonword errors for the repaired completed speech errors is identical with this distribution for undetected completed speech errors (Fisher's exact test: $p = .93$). This suggests that in monitoring overt speech no criterion of lexicality is applied.

3.2 Hypothesis 2. *Overt repairs following interrupted speech errors have shorter offset-to-repair times than overt repairs following completed speech errors.*

As there are no speech wave forms available for the speech errors made in spontaneous speech, this hypothesis will only be tested with the corpus of elicited speech errors.

3.2.1 *Elicited speech errors*

The hypothesis was tested on repaired interruptions and repaired completed exchanges elicited in the experiment described in Nooteboom (2005b) and Experiment 1 as described in Nooteboom and Quené (2008). Although it is assumed that all interruptions have been detected in inner speech, it is not necessarily so that all experimentally elicited interruptions are also repaired, although most are. Probably due to the time pressure in SLIP experiments, of all early interruptions 13% are left without repair and 87% are repaired. Because the dependent measure of interest here is the offset-to-repair time, interruptions not followed by a repair are left out. Only those interruptions were included in which the spoken fragment retained at least the greater part of the vowel. This was done because the same set of cases was used to test the hypotheses on intensity difference and pitch difference between error and repair (see below). The greater part of the vowel was needed for this comparison in intensity and pitch. A "repair" was in most cases identical with the correct target, but in some cases it was not. It was required, however, that the vowels of the spoken error fragment and of the first syllable of the "repair" were identical. In some included cases the interrupted "error" was identical with the beginning of the correct target, and then followed by the correct target as in *kee...keel taart*. There remained 108 interruptions in the two experiments combined that were suitable for inclusion.

For the completed speech errors only those "repaired" errors were included in which the error consisted of a full CVCCVC form, which was either the correct target or the elicited consonant exchange or yet another CVCCVC form. It was required that "error" and "repair" had the same vowel in the first syllable. Quite a few cases were included in which the correct target was repeated. No indication was found that such repetitions differed in any way from real errors followed by the correct targets as repairs. Most of these repetition cases were preceded by an early interruption, as in *gfeit goud...feit goud*. There remained 90 repaired completed errors. The average offset-to-repair time for the 108 repaired interruptions is 182 ms, ranging from 0 to 944 ms, with a standard deviation

of 199 ms, the average offset-to-repair time for the 90 repaired completed speech errors is 398 ms, ranging from 0 to 1371 ms, with a standard deviation of 285 ms. The two means, 182 ms for interruptions and 398 ms for completed errors, differ significantly on a Student's *t* test for independent means ($t = -6.25$; $sd = 242$; $df = 196$; $p < .0001$). These distributions of offset-to-repair times confirm that on the average repairs following early interruptions come much faster than repairs following completed exchanges.

It should be noted that an offset-to-repair time of 0 ms in the case of a completed exchange such as *bood gear* instead of *good beer* corresponds to a much later reaction to the error than an offset-to-repair time of 0 ms in case of an early interruption such as *boo..good beer*, simply because of the time it takes to pronounce the completed exchange. For the CVCCVC word pairs used in this experiment the difference is virtually always in the order of 500 ms. The estimated average difference between the two classes of speech errors in the time needed to react to the overt error therefore is in the order of $(398+500)-182 \approx 716$ ms, and the range of times available for the subject to detect the error and plan a repair runs from 500–1870 ms. Obviously in most completed errors the speakers had plenty of time to detect the error in overt speech and plan a repair. Even in the four cases in which the measured offset-to-repair time indeed was 0 ms, the speakers had some 500 ms during the pronunciation of the error for detecting the error in the first consonant and planning a repair. It has recently been shown that the pronunciation of the error form and planning a repair often go on in parallel (Hartsuiker, Catchpole, De Jong & Pickering 2008). Clearly, on the average repairs of errors detected in inner speech come fast, repairs of errors detected in overt speech come very much slower, probably not only because this time is needed for error detection, but also because the speaker needs this time to plan the repair.

3.3 Hypothesis 3. *Overt repairs following interrupted speech errors have more intensity and higher pitch than the errors themselves. Overt repairs following completed speech errors have less intensity and lower pitch than the errors themselves.*

As no speech wave forms are available for speech errors made in spontaneous Dutch, this hypothesis was only tested with elicited speech errors.

3.3.1 Elicited speech errors

This hypothesis was tested with the same set of cases used for measuring the offset-to-repair times above. In all interruption-plus-repair combinations and all completed-exchange-plus-repair combinations, using the standard intensity display in the PRAAT speech analysis software, the maximum of intensity both in the (first) vowel of the error and in the first vowel of the repair was assessed in decibels SPL, and then the latter was subtracted from the first. Similarly, the pitch in the (first) vowel of the error

and the first vowel of the repair, in both cases at the moment of maximum intensity, was assessed in Hz by taking the inverse of the pitch period duration in seconds. Here also the value found in the repair was subtracted from the value found in the error. The hypothesis states that for interruptions both intensity and pitch have higher values in the repair than in the error, whereas for completed exchanges, both intensity and pitch have lower values in the repair than in the error. The pattern in the data is displayed in Figure 1.




Figure 1. Average difference in intensity (left) and pitch (right) between error and following repair, separately for interruptions and completed exchanges

The average difference in intensity is -1.56 dB ($sd = 3.5$ dB) for interruptions and 2.04 dB ($sd = 4.36$ dB) for completed exchanges, suggesting that indeed for interruptions on the average intensity is higher in the repair than in the error, and for completed exchanges the average intensity is lower in the repair than in the error. The difference between -1.56 dB for interruptions and 2.04 dB for completed exchanges is significant on a student's t test for independent means ($t = 6.46$; $sd = 3.92$; $df = 196$; $p < .001$), confirming that indeed the speakers, after having interrupted speech errors, mostly made repairs with more intensity than the errors, and after having completed exchanges they made repairs with less intensity than the errors.

The average difference in pitch at the moment of highest intensity was -12.93 Hz ($sd = 1.713$ Hz) for interruptions, and 8.62 Hz ($sd = 2.07$ Hz) for completed exchanges. The difference between -12.93 Hz for interruptions and 8.62 Hz for completed exchanges is significantly different on a student's t test for independent means ($t = 8.09$; $sd = 18.7$; $df = 196$; $p < .0001$). Apparently, speakers generally speak the repair on a higher pitch than the error when they are repairing an interruption, and on a lower pitch than the error when they are repairing a completed speech error.

The classification of speech errors as repaired interruptions and repaired completed speech errors most likely does not fully match with the distinction between repairs planned in inner speech and repairs planned in overt speech. Possibly, repairs of interruptions with long offset-to-repair times are only planned during overt speech, and

repairs of completed errors with very short offset-to-repair times were already planned during inner speech. If so, one expects that on the average the intensity difference and pitch difference is a function of the offset-to-repair time: the contribution of misclassified errors becomes bigger when offset-to-repair times get longer for interruptions and shorter for completed errors. Therefore differences in intensity and pitch were correlated with the offset-to-repair times, using Pearson product moment correlations. This was done separately for interruptions and completed speech errors. No significant correlation was found for the interruptions (intensity difference: $r = -.147$; $t = -1.53$; $df = 106$; $p > .1$; pitch difference $r = .0999$; $t = 1.03$; $df = 106$; $p > .3$). However, for the repaired completed errors there are significant positive correlations for both intensity difference ($r = .286$; $t = 2.802$; $df = 88$; $p < .007$) and pitch difference ($r = .3$; $t = 2.999$; $df = 88$; $p < .004$). These correlations suggest that both intensity and pitch in the repair tend to become lowered more relative to the error, as the offset-to-repair time increases. This suggests that repairs of completed speech errors with very short offset-to-repair times more often than repairs with long offset-to-repair times are reactions to inner speech and thus suffer from the tendency to have raised instead of lowered intensity and pitch.

3.4 Hypothesis 4. *Overt repairs following interrupted speech errors are much less often accompanied by editing expressions than overt repairs following completed speech errors.*

This hypothesis will first be tested with elicited speech errors and then with speech errors from spontaneous speech.

3.4.1 *Elicited speech errors*

As it happens, in the repairs of elicited speech errors editing expressions are on the whole not very frequent. In order to have as many cases as possible to test the hypothesis statistically, all repaired interrupted errors and all repaired completed CVCCVC speech errors taken from three experiments, viz. the experiment described in Nooteboom (2005b) and both experiments described in Nooteboom and Quené (2008), were included, not only the many completed and interrupted exchanges in the test condition, but also the much less frequent completed and interrupted exchanges in the base-line condition, i.e. the condition where no segmental biasing was used to elicit exchanges. The relevant pattern in the data is given in Table 3.

Table 3. Repaired interrupted and repaired completed CVCCVC speech errors taken from three experiments (see text), separately for repairs accompanied by an editing expression and repairs not accompanied by an editing expression.

	interruptions	completed exchanges
editing expression	15	15
no editing expression	230	32

In these three experiments combined and the test and base-line conditions combined there are 245 interrupted speech errors. Only 15 (6%) of these are accompanied by an editing expression. Editing expressions are *er*, *hhh*, *sorry*, *of zo iets (or something like it)*, *no*, *mmm*, *hm*, *umnja*, *kweenie (dunno)*, *geen idee (no idea)*. In the three experiments and the test and base-line conditions combined there are 221 completed exchanges. Of these only 47 are followed by a repair. Very likely, most of the other 174 exchanges were not detected by the speaker.

Of the 47 repairs of completed exchanges, 15 (32%) are accompanied by an editing expression. Editing expressions are *ja (yes)*, *er*, *pffhum*, *mm*, *hi*, *andersom (the other way round)*, *ahh...sorry*, *oh*, *hahaha*, *correctie (correction)*, *hm*. The distribution of the numbers of detected speech errors with and without an editing expression is significantly different for interruptions and completed exchanges (Fisher's exact test: $p < .0001$). As predicted, editing expressions are rare for repairs of interruptions and relatively frequent for repairs of completed exchanges. This supports the idea that speakers have a tendency to draw the attention of their audience away from the observable consequences of speech errors detected in inner speech, but tend to direct the attention of their audience to the speech errors they have detected in their overt speech.

Following the same reasoning as applied earlier to the differences in intensity and pitch between error and repair, one would predict that the probability of a repair being accompanied by an editing expression increases with increasing offset-to-repair time. It turns out, however, that far too few editing expressions are made to test this hypothesis statistically.

3.4.2 Speech errors in spontaneous speech

Numbers of repaired segmental speech errors with and without an editing expression were counted in the Utrecht corpus of speech errors, separately for early interrupted speech errors and completed speech errors. Early interrupted here means that at most the initial CV of the error form was pronounced before interruption. Completed speech errors were all those cases where the error form was completed, and the speaker stopped only then for making a repair. It was found that in repairing segmental speech errors speakers hardly ever spoke more than the error form before stopping. They also included hardly ever more than the correct target form itself in the repair (See Nooteboom 2005a: 174–175). There are 122 repaired early interruptions in the collection. The relevant data are given in Table 4.

Table 4. Repaired interrupted and repaired completed segmental speech errors taken from a corpus of speech errors made in spontaneous Dutch (see text), separately for repairs accompanied by an editing expression and repairs not accompanied by an editing expression

	interruptions	completed exchanges
editing expression	5	38
no editing expression	117	210

Only 5 (4%) of the 122 repairs following interrupted speech errors are accompanied by an editing expression. Thirty-eight (15%) of the 248 repairs following completed speech errors are accompanied by an editing expression. The distribution of overt repairs with and overt repairs without an editing expression is significantly different for early interruptions and completed speech errors (Fisher's exact test: $p < .002$). Obviously repairs of interrupted errors are only rarely accompanied by editing expressions and repairs of completed errors are more often accompanied by editing expressions. This confirms that speakers strive towards making a repair as inconspicuous as possible for their audience when they have detected an error in inner speech, but have a tendency of making a repair conspicuous for the listeners when they have detected an error in their overt speech. It should be noted, though, that speech errors made in spontaneous speech seem even less often accompanied by editing expressions than experimentally elicited speech errors.

4. Discussion

In the introduction to this paper it was argued that monitoring for speech errors in inner speech potentially differs from monitoring for speech errors in overt speech, because these two forms of self-monitoring have different functions. When an error is detected in inner speech the speaker strives towards preventing this error from becoming public, preferably with as few perceivable consequences as possible. Because there is very little time to do this, monitoring for speech errors in inner speech is under time pressure, whereas monitoring for errors in overt speech is not. One consequence of this is, as has been confirmed in the present analysis, that, other things being equal, more nonword than real-word segmental errors are detected and then early interrupted, but that the probability of being detected and repaired is the same for nonword and real-word completed segmental errors. The reader may recall that this state of affairs was predicted in the introduction to this paper from the earlier published finding that under time pressure monitoring for speech errors detects more nonword than real-word errors, but when the time pressure is removed, this is not so (Nooteboom & Quené 2008), together with the above-mentioned assumption that monitoring inner speech operates under time pressure, but monitoring overt speech does not. It should be pointed out, however, that the current data have confirmed something that is not easily explained theoretically. Recently, Nozari and Dell (2009) have gone a long way, both in arguing and in demonstrating experimentally that a "lexical editor", being part of a monitor for speech errors, makes little sense, because (1) a lexical editor is too slow to be of much use, and (2) it adds little or nothing to an editor that compares each potential error with the correct target. These authors assume that the working of a "lexical editor" as part of the monitor is similar to a "lexical decision task". In a lexical

decision task participants push one of two buttons, *word* or *nonword*. Nozari and Dell point out that, other things being equal, in a lexical decision task reaction times are much longer than in a task in which participants push one or two buttons, to indicate whether a stimulus word is the same as or different from a pre-specified word target (cf. Foss & Swinney 1973). The problem here is (1) that in theories about monitoring for speech errors, the assumption of a direct comparison between error and target is needed anyway, to explain the frequent and rapid detection of real-word segmental errors, and (2) that currently the lexical decision task is the only model we have of a monitor that detects more nonword than real-word errors. The earlier finding that more nonword than real-word errors are detected in inner speech under time pressure but not under relaxed conditions (Nooteboom & Quené 2008), and the current finding that more nonword than real-word errors are detected in inner speech but not in overt speech, for the time being seems theoretically awkward. It is as yet unclear what monitoring mechanism would detect both real-word errors and nonword errors, but under time pressure more nonword errors than real-word errors, and without time pressure equally many.

The idea underlying this paper is that there are two different “strategies” in monitoring for speech errors, one strategy for inner and one for overt speech. This terminology suggests that in both situations some measure of flexible attentional control is involved. The validity of this assumption is possibly not self-evident to all readers. The reader might perhaps have thought that, although possibly monitoring for speech errors in overt speech is semi-conscious and under flexible attentional control, monitoring for speech errors in inner speech surely is fast and automatic and cannot be under attentional control. However, there are quite many published experiments demonstrating that, although the participants may not be aware of this, some measure of attentional control is also exerted over monitoring for speech errors in inner speech. Examples are Baars et al. (1975) and Hartsuiker, Corley and Martensen (2005). Both papers show that pre-articulatory editing is sensitive to the lexicality of the context. Motley (1980) demonstrated that the probability of a consonant exchange increases considerably when the error is preceded by word pairs creating a semantic bias for the error, such as *bad mug* turning into *mad bug*, when preceding word pairs included *angry insect* and *irate wasp*. Motley (1980) also showed that a conversational setting affected the probability of specific errors: *shad bock* turning into *bad shock* became more likely when the participants expected an electrical shock, and *goxi firl* turning into *foxy girl* was stimulated by an attractive and provocatively attired experimenter. Another example is provided by Motley et al. (1982), who found an effect of the taboo status on the probability of error detection in inner speech. Also the finding by Nooteboom and Quené (2008) that time pressure affects the relative probabilities of error detection suggests some measure of attentional control. All this is relevant, because the current attempt to show that there are different behavioural patterns in repairing errors detected in

inner speech and repairing errors detected in overt speech, derives its logic precisely from the assumption that to a certain extent these different strategies are strategic, stemming from communicative needs. It should also be noted that the actual patterns found, seem to make sense in this context.

If indeed speakers strive to hide, or make less noticeable, for their listeners the perceivable consequences of errors detected in inner speech, they are well advised to keep the offset-to-repair times short. This is precisely what they often seem to do, much more often than when they repair errors detected in overt speech. Admittedly, there are also cases of repairs of interrupted speech errors with longish offset-to-repair times. But obviously, precisely in studying elicited errors of speech and their repairs, one is dealing with an attentional control system that is overburdened. The detection of an error in inner speech, and the decision to interrupt the error as soon as possible, compete under time pressure with the control needed to come up with the correct target as a repair. Interestingly, there are quite a few cases in which the speaker after an interruption hesitates, and then repeats the error that apparently still internally competes with the correct target.

If the speaker (or some subconscious part of the speaker's mind) wants to distract the attention of listeners from the error fragment that was just inadvertently spoken, it may be a good idea to speak the following repair not only fast, but also both louder and on a higher pitch than the error fragment itself was spoken. This is precisely what speakers do. This does not mean that the first syllable of the repair is getting a pitch accent, where the error fragment had not. Virtually all participants in the speech error elicitation experiments speak all regular responses with something that might best be described as double-stress. The impression is definitely that speaking the repair louder and higher-pitched than the error, holds for both syllables, and is not a function of making a pitch accent. In this context it may be relevant that Shattuck-Hufnagel and Cutler (1999) showed results on taped speech errors in spontaneous speech, suggesting that repairs of segmental errors are not marked with an (extra) pitch accent, whereas repairs of lexical errors, involving morphemes or whole words as displaced units, are more often marked with an extra pitch accent.

If the speaker has made a full error, that can hardly have been missed by the audience, he or she should not attempt to hide the fact that an error has been made, but instead should focus the attention of the audience on both the error and the repair. As the error has already been spoken, the only way to do this prosodically is by adapting the prosody of the repair. This can most easily be done by deviating from the often repeated sing song prosody of the regular responses, that of course was also given to the error (because when the error was made it was not yet detected as error). Speaking with somewhat less intensity and on a somewhat lower pitch is a most economical way to deviate from the regular pattern. And this is what the speakers generally do. The lowered pitch makes the impression of some form of finality, as if the repair is more of a final response than the error was. This seems communicatively effective under the circumstances.

Unfortunately, it must remain unclear here how speakers treat their repairs of interrupted and completed speech errors prosodically in spontaneous speech. Shattuck-Hufnagel and Cutler (1999), in investigating acoustically the accent patterns in repairs of segmental and lexical speech errors, did not make a distinction between interruptions and completed errors. This is something awaiting further research on a collection of taped speech errors in spontaneous speech. With respect to the use of editing expressions like *sorry*, *hahaha*, *wrong*, *no*, *er*, etc. speakers of both experimentally elicited speech errors and speech errors in spontaneous speech behave in a communicatively effective way. They add rarely editing expressions to repairs of errors detected in inner speech and more often to repairs of errors detected in overt speech. Levelt (1983) found, in a task where participants had to describe simple networks, that 62% of all error repairs (including not only repairs of segmental errors but also of errors involving lexical units) were accompanied by editing expressions. This is much more than found in the current analysis. Probably, in the SLIP experiments the use of editing expressions is constrained by the time pressure the participants are under. In the corpus of spontaneous speech editing expressions may be rare because they are not always observed or noted down by the collectors. Also in this respect a re-analysis of existing corpora of taped speech errors in spontaneous speech would be welcome. The scarcity of editing expressions in the current data also prohibits further analysis of the possible differences in the use of specific editing expressions, although it seems to be case that most editing expressions tend to follow the repair, but that *uhh* or *er* more frequently precedes the repair. This confirms that *uhh* or *er* first of all serves to hold the floor while a repair is being planned (cf. Levelt 1989: Chapter 12).

The current analysis supports an important aspect of Levelt's perceptual loop theory of monitoring, viz. that both inner and overt speech are being monitored for speech errors (see also Hartsuiker & Kolk 2001; Hartsuiker, Kolk & Martensen 2005). There is at least a delay of 200 or 250 ms between the two (Hartsuiker & Kolk 2001), and potentially much more, depending on how much material is buffered in inner speech. If an error is not detected in inner speech, detection of this error in overt speech is of course perfectly appropriate. However, it is unlikely that speakers can monitor inner and overt speech with equal attention at each moment in time. It has been assumed that monitoring for speech errors is under attentional control (Hartsuiker, Kolk & Martensen 2005) so that the speaker may either direct his or her attention more to inner speech or more to overt speech. In this respect it is revealing that the current analysis shows that in SLIP experiments there are far more interrupted errors (169) than repaired completed errors (29). Apparently, during these experiments the participants' attention is focused on inner speech. But the data on speech errors in spontaneous speech show that there are more repaired completed speech errors (244) than interrupted speech errors (179), showing that in normal spontaneous speech monitoring attention is certainly not less for overt speech than for inner speech. This makes sense for two reasons. First, in SLIP experiments participants have little time to detect and repair errors in overt speech.

Immediately after a response, the next series of biasing stimuli starts. Secondly, in SLIP experiments participants do not have to worry about detecting lexical errors, but in spontaneous speech this is an important aspect of monitoring. These lexical errors are probably much harder to detect in inner speech than in overt speech, because the phrasal context is needed for error detection (cf. Nooteboom 2005a).

The current findings complicate our view of self-monitoring. Whereas for a long time it was thought that self-monitoring operates in the same way whether it is directed at inner speech or at overt speech, using a few straightforward criteria for filtering out speech errors (cf. Levelt 1989), it must now be admitted that this view was too simple. The ways self-monitoring for speech errors operates are controlled by different goals in inner and in overt speech, and these different goals lead to different behavioural patterns. It seems indeed the case that speakers attempt to make speech errors detected in inner speech as little conspicuous as possible and to draw the listeners' attention to repairs of speech errors detected in overt speech. This explains the different behavioural patterns found in these two classes of overt repairs of detected speech errors.

References

- Baars, B.J., Motley, M.T. & MacKay D. 1975. Output editing for lexical status from artificially elicited slips of the tongue, *Journal of Verbal Learning and Verbal Behavior* 14: 382–391.
- Blackmer E.R. & Mitton J.L. 1991. Theories of monitoring and the timing of repairs in spontaneous speech. *Cognition* 39: 173–194.
- Del Viso, S., Igoa, J.M. & Garcia-Albea, J.E. 1991. On the autonomy of segmental encoding: Evidence from slips of the tongue in Spanish. *Journal of Psycholinguistic Research* 20: 161–185.
- Dell, G.S. 1986. A spreading-activation theory of retrieval in sentence production. *Psychological Review* 93: 283–321.
- Dell, G.S. & Reich, P.A. 1981. Stages in sentence production: An analysis of speech error data. *Journal of Verbal Learning and Verbal Behavior* 20: 611–629.
- Foss, D. & Swinney, D. 1973. On the psychological reality of the phoneme: Perception, identification, and consciousness. *Journal of Verbal Learning and Verbal Behavior* 12: 246–257.
- Garrett, M.F. 1976. Syntactic process in sentence production. In *New Approaches to Language Mechanisms*, Wales, R.J. & Walker, E.C.T. (eds), 231–256. Amsterdam: North-Holland.
- Hartsuiker, R.J. 2006. Are speech error patterns affected by a monitoring bias? *Language and Cognitive Processes* 21(7–8): 856–891.
- Hartsuiker R.J., Antón-Méndez I., Roelstraete B. & Costa A. 2006. Spoonish spanerisms: A lexical bias effect in Spanish. *Journal of Experimental Psychology, Learning, Memory and Cognition* 32 (4): 949–53.
- Hartsuiker, R.J., Catchpole, C.M., De Jong, N.H. & Pickering, M.J. 2008. Concurrent processing of words and their replacements during speech. *Cognition* 108(3): 601–607.
- Hartsuiker, R., Corley, M. & Martensen, H. 2005. The lexical bias effect is modulated by context, but the standard monitoring account doesn't fly: Related reply to Baars, Motley, & MacKay (1975). *Journal of Memory and Language* 52: 58–70.
- Hartsuiker, R.J. & Kolk, H.J. 2001. Error monitoring in speech production: A computational test of the perceptual loop theory. *Cognitive Psychology* 42: 113–157.

- Hartsuiker, R.J., Kolk, H.J. & Martensen, H. 2005. Division of labor between internal and external speech monitoring. In *Segmental Encoding and Monitoring in Normal and Pathological Speech*, R. Hartsuiker, Y. Bastiaanse, A. Postma & F. Wijnen (eds), 187–205. Hove: Psychology Press.
- Kolk, H.J. 1995. A time-based approach to agrammatic production. *Brain and Language* 50: 282–303.
- Levelt, W.J.M. 1983. Monitoring and self-repair in speech. *Cognition* 14: 41–104.
- Levelt, W.J.M. 1989. *Speaking: From Intention to Articulation*. Cambridge MA: The MIT Press.
- Levelt, W.J.M., Roelofs, A. & Meyer, A.S. 1999. A theory of lexical access in speech production. *Behavioral and Brain Sciences* 22: 1–75.
- Motley, M.T. 1980. Verification of ‘Freudian slips’ and semantic prearticulatory editing via laboratory-induced spoonerisms. In *Errors in Linguistic Performance: Slips of the Tongue, Ear, Pen, and Hand*, V.A. Fromkin (ed.), 133–148. New York NY: Academic Press.
- Motley, M.T., Camden, C.T. & Baars, B.J. 1982. Covert formulation and editing of anomalies in speech production: Evidence from experimentally elicited slips of the tongue. *Journal of Verbal Learning and Verbal Behavior* 21: 578–594.
- Nooteboom, S.G. 1980. Speaking and unspeaking: Detection and correction of segmental and lexical errors in spontaneous speech. In *Errors in Linguistic Performance: Slips of the Tongue, Ear, Pen, and Hand*, V.A. Fromkin (ed), 87–95. New York NY: Academic Press.
- Nooteboom, S.G. 2005a. Listening to one-self: Monitoring speech production. In *Segmental Encoding and Monitoring in Normal and Pathological Speech*, R. Hartsuiker, Y. Bastiaanse, A. Postma & F. Wijnen (eds), 167–186. Hove: Psychology Press.
- Nooteboom, S.G. 2005b. Lexical bias revisited: Detecting, rejecting and repairing speech errors in inner speech. *Speech Communication* 47(1–2): 43–58.
- Nooteboom, S.G. 2007. Alphabetics. From phonemes to letters or from letters to phonemes? *Written Language and Literacy* 10(2): 129–144.
- Nooteboom, S.G. & Quené, H. 2008. Self-monitoring and feedback: A new attempt to find the main cause of lexical bias in segmental speech errors. *Journal of Memory and Language* 58: 837–961.
- Nozari, N. & Dell, G.S. 2009. More on lexical bias. How efficient can a “lexical editor” be? *Journal of Memory and Language* 60(2): 291–307.
- Shattuck-Hufnagel, S. & Cutler, A. 1999. The prosody of speech error corrections revisited. Proceedings of the International Congress of Phonetic Sciences, August 1–6, 1999, 1483–1486, in San Francisco. San Francisco: Regents of the University of California. (on CD-ROM).
- Stemberger, J.P. 1985. An interactive activation model of language production. In *Progress in the Psychology of Language*, Vol. 1, A.W. Ellis (ed), 143–186. Hillsdale NJ: Lawrence Erlbaum.

What's in a quantifier?

Rick Nouwen*
Utrecht University

In this article, I discuss several inquiries into the meaning of expressions of quantity. It is to some extent received wisdom that quantifiers are not expressions that stand proxy for amounts, quantities, frequencies, etc. In cognitive psychology, a common observation is that quantifiers express perspectives on quantity. I will argue that while this result bears some relation to notions known from formal semantics, the dominant theoretical framework, generalised quantifier theory, nevertheless falls short of providing a unifying approach to logical, linguistic and psychological aspects of quantity expressions. The upshot is that the many diverse ways in which quantity information may be worded should be reflected in a variety of analyses. In other words, quantifiers need to be approached on a case by case basis.

1. Introduction

The democracy of the internet can appear to give rise to some amazing insights. Jyte.com is a website where one can submit statements which are evaluated by visitors of the site who vote either in favour or against. After someone submitted the claim in (1), a striking discussion about quantifiers emerged.

- (1) I tend to take *a couple* to mean exactly two.

In Dutch, my native language, *een paar* (literally, a pair) is, at first sight, used in much the same way as *a couple* is in English. However, it is perfectly fine for me to say (2) in a situation where I got three or four books out of the library.

- (2) Ik heb een paar boeken uit de bibliotheek gehaald.
I have a pair books from the library taken
'I got a few books from the library.'

My personal experience with (British) English is that *a couple* tends to be somewhat more faithful to the lexical meaning of *couple*: *to invite a couple of friends* means to

*This work is supported by a grant of the Netherlands Organisation for Scientific Research (NWO), which is hereby gratefully acknowledged.

invite two friends, no more. The democracy of jyte.com, however, proved my intuitions partially wrong: eighteen voters agreed with (1), sixteen disagreed. Several reasons were given for disagreeing, ranging from plain intuition to arguing that *a couple* can't mean *two* because *two* already does.

The statement in (1), however, sparked the interest of the jyte community into quantifiers, resulting in a number of statements designed to try and tackle the question of how quantifiers compare to one another. For instance, there was widespread agreement (34 in favour 0 against) that *a pair* means exactly two, and that *a handful* is more than *a couple* (27–0). Here is an overview of the resulting votes with respect to some other quantifiers.

(3)	lots > many	16–7
	many > several	23–1
	several > a few	24–1
	a few > a handful	2–21
	oodles > lots	24–2

Such comparisons reflect a strange frustration language users have about the expression of quantity, namely that there appears to be no neat order of things. We never acquire any specific knowledge of how many *many* really is, nor what the difference (in quantity) is between *many* and *lots*.

This frustration has also been a scientific one. As I will discuss in Section 2, the jyte discussion is reminiscent of a certain line of psychological investigation that has proven popular in the past decades. From a more theoretical perspective, the meaning of quantifiers has been on the scientific agenda since Aristotle's *de interpretatione*. However, the bulk of the post-Aristotelian discussion will do little to settle issues such as those in (1) and (3). This is because if anything is clear about quantifiers, it is that they are not simply words that stand proxy for amounts, quantities, numbers, frequencies, proportions, etc. That is, as we will see, you do not get to know what a quantifier means by comparing it to other quantifiers or to a (precise) quantity. I will illustrate this from the point of view of theoretical semantics in Section 4, after approaching it from a more psychological perspective in Section 3.

In a nutshell, the problem is that quantifiers form an enormously varied group of linguistic expressions. Some of these are vague, some are imprecise and yet others denote exact quantities. Comparing quantifiers of all kinds with each other in order to find out something about their meaning would miss the point that there are significant differences in how precise, imprecise and vague expressions are used. Many quantifiers moreover encode an evaluation of the quantity they communicate. That is, they approach quantity information from a particular perspective. Such intricacies complicate matters even further.

If, however, comparisons like those of (3) do not help in discovering the meaning of quantifiers, then what will? Below, I will be on the look out for a general model of

quantifier meaning. In particular, I will discuss to what extent a reconciliation between logical, linguistic and psychological approaches is possible and desirable. As we will see, there is a gap between empirically motivated psychological models of quantifier meaning and the models provided by formal semantic theory. This gap will turn out to be difficult to bridge. The modest objective of this paper is to at least pinpoint some of the underlying problems. One of the most serious problems for reaching a satisfying model of quantifier meaning, I will argue, is that there may not really be a way of making theoretical sense of the notion of a *quantifier* itself.

2. Words and numbers

2.1 Ordering quantity expressions

There exists a long line of (largely psychological) literature which questions how people use quantity expressions and how, in language use, such expressions line up (Simpson 1944 is the earliest attempt I know). The *ordering* of expressions is often thought to be important for the purpose of understanding the communication of quantitative information, especially in domains where more precise information is valuable, such as is the case for conveying medical information. O'Brien (1989) asked general practitioners to imagine a situation where they had to communicate the likelihood of there occurring a certain side-effect with a drug. They then had to map twenty-three frequency quantifiers onto a percentage scale. In similar experiments also geared at medical practice, Renooij and Witteman (1999) set out to find "an acceptable representation of a probability scale that [contains] mutually exchangeable verbal and numerical expression" (Renooij & Witteman 1999: 30). In other words, the goal of such work is to eliminate the gap between words and numbers, by experimentally finding out how they correspond to one another.

In studies of questionnaire design, quantifiers have received similar attention. Bradburn and Miles (1979) ask how to interpret responses to survey questions made using vague quantifiers. They report on an experiment where subjects fill in a survey where the response categories form a scale of vague (frequency) quantifiers. They are then subsequently asked what their responses meant. That is, somebody answering *often* to a certain question had to specify 'exactly how many times a day' this corresponded to (Bradburn & Miles 1979: 95, see Pace & Friedlander 1982 for a similar study). Pracejus et al. (2004) investigate the interpretation of cause-related marketing slogans that come with products ('a portion of the proceeds goes to charity'). One of their studies involved subjects estimating amounts donated to charity for a purchased item given the product's price and a particular slogan.

A common result of such studies into practical applications of quantifier meaning is a considerable between-subject variation with respect to how a certain quantifier

relates to some value. In many cases, this finding seems to be a major cause of concern for the authors. Pace and Friedlander, for instance, remark that “if there are substantial differences in the way people interpret the labels [i.e. the vague quantifier response categories (RN)], then, unlike standardized objective test data, interinstitutional comparisons are of questionable value” (Pace & Friedlander 1982: 267). Such worries lead researchers to investigate what contextual parameters are responsible for differences in how vague labels are estimated to correspond to numerical information. Schaeffer (1991) considers social factors that might play a role, concluding that “there are significant differences by race, education, and age in the meaning of phrases conveying relative frequency” (Schaeffer 1991: 395).

From such lines of research there seems to spring a general consensus that quantifier-value mappings are hopelessly subject-specific and hardly useful, due to the enormous overlap between different quantifiers. Quite a few researchers saw this as an incentive to investigate the source of the mapping problem, resulting in an ongoing range of mapping experiments, witness some of the more recent papers cited above. One result of O’Brien’s (1989) study of doctor-patient risk communication was that for some expressions there is huge variance in how they map to a scale. Furthermore, he observed that “the context in which a given expression of probability is used may affect its meaning” (O’Brien’s 1989: 98). Nevertheless, O’Brien took the fact that some expressions showed a relatively low amount of variance between subjects to mean that the “results are encouraging and suggest that phrases denoting likelihood might be systematically codified to enhance communication between doctor and patient” (O’Brien’s 1989: 98). This was wishful thinking, rather than an informed estimation of future progress and I contest that in the 30 years after O’Brien’s study we are any nearer to such a systematic codification. In fact, the style of comparison of quantifiers in such studies entails a distorted understanding of what quantifiers do. In the next section, I provide a detailed discussion of a general problem with the above view. In the remainder of the current section, I will briefly reflect on the topic of context-dependence and argue that it is an oversimplification to believe that the problem with the above experiments is simply one that can be solved by fixing the context.

2.2 Context-dependence and non-extensionality

From linguistic semantic theories of quantification it is well-known that some natural language quantifiers are severely context dependent. Compare (4a) to (4b).

- (4) a. Many students came to my party.
- b. Some students came to my party.

Now imagine two different situations. In both, out of the eighty I invited, forty-five students came to my party. In the first situation, I sent invitations indicating that there

would be an entrance fee to the party of €60 and that people had to bring their own drinks. In the second, the invitations mentioned free beer, free food and no entrance fee. Although in both situations, the number of students coming to my party is forty-five, (4a) is much more likely to be true in the first situation, than in the last. We expect few students to come to expensive parties and many to come to parties with free food and beer. So, a forty-five out of eighty score counts as *many* with respect to the first situation, but probably not with respect to the second. This contrasts with (4b), which seems true irrespective of the situation. In other words, *many*, but not *some* is context-dependent.

Context-dependence also surfaces when we compare different sentences with the same quantifier.

- (5) a. There are many ants in my garden.
- b. There are many typos in this article.

The average garden has millions of ants, while the average manuscript probably has around twenty typos. So, when someone's garden counts one hundred ants, this person is unlikely to utter (5a). However, a manuscript with one hundred typos is likely to be described as in (5b). This is context-dependence, again. *Many*, when talking about ants differs from *many* when talking about typos.

But things are even more complicated. Keenan and Stavi (1986) and Westerstahl (1984) show that contrasts like that in (5) surface even when *many* talks about the same group of individuals, in the same context. Imagine a conference of lawyers and policemen where normally 60 lawyers and 40 policemen attend. Also, on average, only 10 attendants are women. This year, there are only 20 lawyers, but a staggering 80 policemen. Strikingly, all the lawyers happen to be women and all the policemen are men. In this (bizarre) situation (6a) is probably judged to be false, but (6b) tends to be judged as true.

- (6) a. Many lawyers attended the meeting this year.
- b. Many women attended the meeting this year.

The thing to realise is that the set of attending women is the set of attending lawyers. This shows that if the context and the number of relevant individuals are both fixed, *many* still gives rise to different meanings. Counting women in the above situation leads to the same amount as counting lawyers, but because the expectancy for lawyers and women differ, *many* is sensitive to what is counted.¹

1. This example is from Keenan and Stavi 1986. Westerstahl 1984 shows a similar effect for when the verb phrase is varied:

- (i) a. Many students in the class got the highest mark on the exam.
- b. Many students in the class are right-handed.

Keenan and Stavi and Westerståhl conclude from these observations that *many* is *non-extensional*. This means that whether *many* lawyers attended a meeting cannot be decided on the basis of looking at the actual set of lawyers at the meeting (in relation to the actual full set of lawyers) only. Lappin (2000) proposes to account for this intuition by assigning a semantics to *many* that crucially depends on normative situations.

- (7) Lappin (2000) (paraphrased): *Many A's B* is true if and only if for every normative situation it is the case that the actual number of A's that B exceeds the number of A's that B in the normative situation.

So, if in a normal situation, 60 to 80 lawyers turn up, *many lawyers turned up* will mean that more than eighty of them came. Likewise if in a normal situation, 20 to 30 women turn up, *many women turned up* will mean that more than thirty women came. This difference arises from what is normal and is independent of the fact that in the actual situation the number of lawyers and women is the same.

Presumably, one could, and should, understand the notion of *normative situation* broadly, even containing notions like *expected situations*, desired situations, etc. For example, (8a) and (8b) do not necessarily depend on what the average situation looks like.

- (8) a. Many people saw through my disguise. *More than I expected*
 b. There are many people in this queue. *More than I want there to be*

What the above shows is that *many* is not just a word that stands in for an amount. It has a complex non-extensional relation with an amount, which depends on what is normality, what is expected, what is desired, etc. If we are optimistic, we could say that there could be a psychological agenda to find out how many is *many* given a context and given the speaker's expectations and such. However, such research would be missing the point that when a speaker uses a vague quantifier like *many*, she may do so with a rather specific purpose. In the literature on vagueness, there is an emerging consensus that vague communication is *useful* (Parikh 1994; Kyburg & Morreau 2000; Barker 2002). One way in which this is manifested is that a speaker may use a vague term, without any intention of communicating the extension of that term. This can be illustrated along the lines of Barker (2002). Take a sentence like *Cindy is tall*. When I say this to someone who does not know Cindy, the hearer will receive the vague information that Cindy is taller than some contextual standard. But when I say this to someone who knows Cindy and moreover who knows exactly how tall she is, the hearer will

Imagine a situation in which in some thirty student class, ten of the students got the highest mark on the exam, but that these ten happen to be all the right-handed students. In such a situation (32a) tends to be thought of as true (one would expect fewer students to have the highest mark), while (32b) tends to be judged false (normally, there would be more right-handed students).

learn something about what I experience to be tall. In other words, in such cases, the information exchange is not about Cindy's height at all. To return to *many*, a sentence like *Many students came to my party* might have as its sole communicative function the message that the speaker is relieved, surprised or shocked by how many students came. In that sense, the sentence is not intended to communicate a quantity at all.

The moral of the above is that the elimination of vagueness from a term, strips this term from some of its key communicative functions. This still leaves open, however, the question of how vague quantifiers compare to each other. If *few* and *a few* are both vague, then what drives the differences in their usage? This is the topic of the next section.

3. Directivity

In one particular line of psychological research into quantity expressions, many arguments have been formed against reducing the meaning of quantity expressions to quantities. One particularly enlightening insight from that area is the idea that asking a subject to place multiple quantifiers on a scale creates a bias. Moxey and Sanford (1993, cf. Moxey & Sanford 2000) gave a large group of subjects the task of assigning a number to a quantifier. Each subject, however, only got a single quantifier, which means that there could be no effect caused by the act of comparing one expression with another. On the basis of the results of this experiment, quantifiers like *few*, *very few*, *only a few*, *not many* and *a few* turned out to be indistinguishable.

Moxey and Sanford took their results to indicate that there is more to quantifier meaning than the indication of amounts, numbers or frequencies. In an influential research programme, they set out to show that quantifiers carry information about the *perspective* that is taken. The idea is that certain quantifiers express the same quantity, yet differ in how they focus on this quantity. For instance, it is conceivable that in a situation where three people came to my party I can express this with both (9a) and (9b).

- (9) a. Few people came to my party.
- b. A few people came to my party.

There is a difference, however. Whereas *few* is negative in flavour, *a few* is not. Below we turn to the question how to make such a flavour precise, but for now it suffices to point out the intuition that (9a) expresses a negative perspective of how many people came, whereas (9b) is positive in nature. There seems to be a variety of ways to refer to the negative/positive distinction, such as polarity, perspective and focus. These terms are rather discipline-specific and often come with a particular view on the distinction. In order to have a more pre-theoretic terminology, I will use the term *directivity*.

Directivity also creates a contrast in (10). In terms of the number of people who came to my party, (10a) and (10b) say the same two things: (i) some did not come, (ii) but those who did not were few.

- (10) a. Not quite all people came to my party.
b. Almost all people came to my party.

Nevertheless, sentences like (10a) are often called negative, while (10b) is positive.² How can you tell? One of the most powerful diagnostics is to mark the sentence with an evaluative predicate, like *fortunately* or *surprisingly*, as in (11).

- (11) a. Surprisingly, not quite all people came to my party.
b. Surprisingly, almost all people came to my party.

These sentences say different things. In (11a), the speaker expected everyone to come. In (11b), however, the speaker expected not so many people to come. Such effects are quite robust, and are certainly not limited to a select group of evaluative adverbs. For instance, Anscombe and Ducrot (1976) notice that something like directivity influences how pieces of information may or may not connect in discourse. For instance, in (12), the information provided by the negative phrase *few motorists* is given a specification in a parenthetical. In (12a) this specification is made using the positive quantifier *almost 20%*, which somehow clashes with the negative perspective provided by *few*.³ (See Jayez & Tovenas 2007 for a formal semantic discussion of such examples.)

- (12) a. ??Few motorists, almost 20%, know the details of the new highway code.
b. Few motorists, about 20%, know the details of the new highway code.

These particular kinds of manifestations of the notion of directivity have received considerable attention in psychology as well. Sanford et al. (2002), for instance, focus on contrasts like that in (13) and experimentally confirmed the intuitions for them.

- (13) a. In the airplane crash, a few people were killed, which is a #good/bad thing.
b. In the airplane crash, few people were killed, which is a good/#bad thing.

On an intuitive level, such contrasts may be explained by the idea that a negative quantifier like *few* in (13a) focuses on the people that were not killed, while the positive quantifier

2. An anonymous reviewer justly complains that the labels *positive* and *negative* are always interpreted differently by different people. As such, they are confusing terms. However, I will use them in the way they are quite often used in the literature, as technical predicates that indicate a theoretically relevant distinction in directivity. This will become clearer from the discussion below.

3. See Horn (2002) for similar contrasts. For instance, (ii) illustrates that different kinds of exemplification may point to a positive/negative contrast (cf. Dayal 1998).

(ii) a. A few students came to my party. For example, Bill did / ??didn't.
b. Few students came to my party. For example, Bill ??did / didn't.

Whereas *a few* triggers plain exemplification, negative *few* triggers counter-exemplification.

focuses on the ones that were. This intuition is supported by a large set of experimental results on how directivity interacts with anaphora.

Quantifiers give rise to various kinds of (plural) discourse anaphora, involving reference to a variety of sets associated with the quantified sentence. In a sentence like *Most students came to my party*, the (contextual) set of students is called the maximal set (or maxset). The set of students that came is called the reference set (or refset). Finally, the set of absent students is called the complement set (or compset). These three sets give rise to three kinds of pronominal anaphoric reference, labelled maxset reference, refset reference and compset reference. They are exemplified in (14).

- (14) a. Most students came to my party. Some of them, however, stayed at home instead.
 them = maxset
- b. Most students came to my party. They had a wonderful time.
 they = refset
- c. Very few students came to my party. They stayed at home instead.
 they = compset

The phenomenon in (14–c) is also termed *complement anaphora*. From a series of experiments conducted in the eighties and nineties (see for instance, Moxey & Sanford 1987; Sanford & Moxey 1993; Paterson et al. 1998 and references therein) it became clear that this latter form of discourse anaphora is associated with quantifiers with a negative perspective. In fact, in continuation experiments where subjects were presented with stimuli like (15) (with Q some quantifier), positive quantifiers like *most* only exceptionally lead to completions where the pronoun refers to the complement set.

- (15) Q of the MPs attended the meeting. They. . .

Pairs like *few* and *a few* show strikingly different results. For instance, in Moxey and Sanford (1987), *few* and *very few* trigger many completions with complement anaphora, whereas *a few* only does so in 5% of the cases. Similar effects are found in the interpretation of inclusion relations (which are, arguably, anaphoric in nature).

- (16) a. Not many of the fans went to the match, including Fred.
 b. A few of the fans went to the match, including Fred.

When asked on the basis of examples like these whether Fred did or did not go the match, subjects tend to say no to (16a) and yes to (16b) (Sanford et al. 2001).

In sum, we have seen a wealth of effects which show that quantifiers have a richer meaning and use than the mere indication of a quantity. Several models have been proposed for these effects, which are generally based on the recognition that negative quantifiers are somehow special. Rather than simply asserting a quantity, these quantifiers make

salient a difference from expectation.⁴ (See Sanford et al. 2007 for a precise and recent exposition.) Such a psychological model advocates the understanding that quantifier meaning is something which is more complex than quantity assertion. In the next section, I turn to the theoretical linguistic semantics and its approach to quantifier meaning, and, subsequently, I will ask whether in that discipline a result similar to that of psychology is possible. In other words, the question is whether the parameters of the psychological model correspond to parameters known in semantic theory.

4. Quantifier semantics

So far, I have been loosely using terms like *quantity expressions*, *quantifiers* and *determiners*. It makes a lot of sense, however, to be quite a bit more explicit, since the quantifier literature is laden with terminological chaos. In the process of doing so, I introduce the single most influential theory on quantifier semantics. Subsequently, I will apply the state of the art in quantifier semantics to the effects of directivity discussed above. I will conclude by questioning the generality of the semantic framework.

4.1 Generalised Quantifier Theory

At least superficially, there is an obvious syntactic similarity between expressions such as *some*, *every*, *many*, *most*, *no*, *more than fifty-five*, *between nine and ninety*, etc. Aristotle already identified part of this class as responsible for creating the subjects that together with predicates form propositions. In somewhat more modern terminology, Aristotle's characterisation translates in the understanding that these expressions form noun phrases, which in combination with verb phrases form sentences. According to the leading view on quantifiers in formal semantics, Generalised Quantifier Theory (Barwise & Cooper 1981; Keenan & Stavi 1986; van Benthem 1986), this syntactic similarity between the expressions is mirrored in semantics: they all express the same kind of semantic entity. In other words, from a semantic point of view, these words and phrases form a coherent class: they are a (special kind of) quantifier.

There are many unfortunate terminological issues in this part of semantics, all to do with the notion of a quantifier. Let us settle some of these first. Many things have been called a quantifier in many different disciplines. In linguistics, there are those who refer to the determiner(like) elements mentioned above as quantifiers (e.g. *every*), and others who reserve this term for determiner phrases (e.g. *every student*). From a mathematical point of view, both are correct in a sense. A mathematical (generalised)

4. That is, in this sense negative quantifiers are reminiscent of Wason's (1959) characterisation of negative statements in terms of deviation from what is expected.

quantifier is namely any relation between subsets of or relations on the domain of entities (Mostowski 1957; Lindstroem 1966). This gives rise to a typology of quantifiers. Unary relations between sets are called type (1) quantifiers, binary relations between sets are of type (1,1). A binary relation between sets and binary relations is of type (1,2), etc.

The application of mathematical generalised quantifiers to linguistics is one of the true success stories of formal semantics. Montague (1973) as well as Lewis (1970) provided the basis for a range of papers in the early eighties which in hindsight are collectively referred to as *generalised quantifier theory* (GQT, a.o. Barwise & Cooper 1981; van Benthem 1986; Keenan & Stavi 1986; Westerståhl 1984). The central idea is that there is a correspondence between noun phrases (or, in modern speak, determiner phrases) and type (1) quantifiers. Moreover, these type (1) quantifiers are the result of filling in one of the arguments of a type (1,1) quantifier. In other words, determiner-like elements correspond to (1,1) quantifiers. In what follows, I reserve the term natural language quantifier (or, if the context is in no need of such an elaborate qualification, simply quantifier) for this latter category.

As an illustration, consider the case of *every*. In GQT, *every* denotes a relation between sets, or properties. Two properties are in the *every* relation if the one is included in the other. So, *every raven is black* is true if the property of being a raven includes the property of being black. In set theoretic terms: the set of ravens is included in the set of black things. *Every raven* is now a collection of properties (i.e. a set of sets): the set of properties that is shared by every raven, which includes things like being black, having wings, having a curved bill, being a bird, etc. Schematically:

expression	Denotation	Type
every	The relation between properties A and B such that every individual that has A has property B	(1,1)
every raven	The set of properties that every raven has	(1)

This kind of semantics for quantifiers helps us to understand what goes wrong when we try to compare expressions of quantity to numbers, proportions or amounts. Take *most*. A fairly accurate semantics (though see Ariel 2004) is given in (17), which says that *most* expresses a majority relation.

- (17) *Most A B* is true if and only if there are more individuals that have both property A and B than there are individuals that have property A but not B. (In other words, more than fifty percent of the individuals that have property A, have property B as well).

According to this, *most* simply gives a lower bound. As soon as more than fifty percent of the A's have property B, then *most* can be truthfully used. In other words, its semantics is

very imprecise in the sense that it covers a wide range of percentages. Moreover, there is a sense in which this semantics predicts *most* to be positive. As said, (17) specifies a lower bound. This means that *most* is compatible with the top of the scale, but not with the bottom. The semantics in (17) predicts *Most A B* to be true, if all A's have property B, but not if none has. This accounts for (18).

- (18) Most students passed the test. In fact, all/#none did.

In fact, there is a more general way of approaching directivity in generalised quantifier theory, namely by applying the formal notion of monotonicity. In general, an operation on a domain is *monotone* if applying the operation preserves a certain structure that exists within this domain. A quantifier is monotone when the quantifier preserves (or reverses) the subset relation between sets.⁵ This is best explained by an example.

The set of people with hair has as one of its subsets the set of people with curls. *Most* is upward monotonic (or monotone increasing) because if I know that *most professors have curls* (the subset case), I automatically know that *most professors have hair* (the superset case). *Less than half* is downward monotonic (or monotone decreasing), for if I know that *less than half the professors have hair* (the superset case) then *less than half the professors have curls* (the subset case). *Exactly half* is non-monotone: if *exactly half the professors have hair*, nothing may be concluded about whether or not *exactly half the professors have curls*, nor is there an inference the other way around.

It can now be hypothesised that negative quantifiers are those that are downward monotone. For instance, *a few* is clearly upward monotone, while *few* appears to be downward oriented. This is in accordance with our perception of *few* as negative and *a few* as positive.

- (19) a. A few professors have curls. Therefore, a few professors have hair.
b. Few professors have hair. Therefore, few professors have curls.

However, there is a problem with the context-sensitivity of a vague quantifier like *few*. One could, for instance, question the validity of (19b) by imagining a situation in which 95% of the professors are bald, but where the remaining 5% are all curly. Surely, such a situation makes the premise true, but the conclusion false. However, it is possible to argue that what happens in such cases is that the criteria for *few* in the first sentence are different from those in the second. So (19b) is then no longer a test for monotonicity, since it contains two different quantifiers.

A different kind of test, one which eliminates context effects, clarifies matters for the case of *few* (and other context-sensitive quantifiers). Say that Bob is a professor

5. To be precise, a type (1,1) quantifier can be monotone with respect both arguments. Monotonicity in the left argument is often called persistence. In what follows, we focus on the second argument only.

and that he has curly hair. If Bob now were to decide to straighten his hair, the set of curly people in the new situation would be a proper subset of the old set of curly people. Intuitively, Bob's sole action does not influence the standard of comparison we associate with quantifiers like *many*, *few* and *a few*, and so by comparing our intuitions about the curly Bob situation (the superset case) to those about the situation after Bob straightened his hair (the subset case), we can eliminate the influence context has on such quantifiers. Such a setup confirms the downward monotonicity of *few*. Clearly, Bob's action will not matter to the truth of *few professors have curls*. Once true for the initial situation, it remains true after Bob's makeover. The same does not apply to *a few*. If at first *a few professors have curls*, Bob's straightening of his hair could render this sentence false, for it could be that Bob was the only curly professor. As already suspected, *a few* is upward monotone.

4.2 Monotonicity and directivity

It should be clear that GQT provides us with a basic yet powerful framework for defining a precise and general quantifier semantics, which comes with a set of tools for identifying properties of quantity expressions. The question is whether these tools suffice for the characterisation of various aspects of quantifier meaning. If negative quantifiers are monotone decreasing quantifiers, then it would be instructive to find an explanation of the diagnostics of negativity in terms of monotonicity.

In Nouwen (2003), I propose that the contrast in (20) can be explained by means of monotonicity. (See Kibble 1997; Hendriks & de Hoop 2001 for related, yet different, proposals.)

- (20) a. Most students saw the match. #They went to a party instead. positive
 b. Few students saw the match. They went to a party instead. negative

The way I argued was that complement anaphora is a marked form of reference, which involves an inference step: the inference of the existence of a (nonempty) compset. With a quantifier like *most* such an inference is not possible.⁶ The reason is that *most A B* is true in case *all A B*. So, on hearing *most students saw the match*, a hearer does not know whether there are any students who did not see the match, and so the existence of a compset can not be inferred. In fact, all monotone increasing quantifiers are compatible with cases in which a universal statement is true and, so, all monotone increasing quantifiers preclude the inference necessary for complement anaphora.

Downward entailing quantifiers will often, but not always, allow for the inference to take place. For instance, if *few* in (20b) is taken to mean something like *less than*

6. A complication is formed by the possibility that *most* might give rise to implicatures. See footnote 12 for discussion.

twenty percent, then we know that at least eighty percent of the students did not see the match. Complement anaphora is therefore licensed. There is a complication, however. Downward monotonicity is not a sufficient condition for inferring a nonempty compset. In a sentence like (21), the modified numeral is downward entailing.

- (21) Less than twenty students saw the match.

Despite downward monotonicity, however, (21) fails to guarantee the non-emptiness of the compset, for this example could be true simply because the total number of students is fewer than twenty. In that case, it could be that there are no students who did not see the match: all of the fewer than twenty students saw the match. Indeed, negative modified numerals are relatively bad licensors of complement anaphora, and so the monotonicity theory of Nouwen (2003) would explain why negativity is a necessary, yet not a sufficient trigger of complement anaphora.

The inferences licensed by monotonicity may also account for the way evaluatives interact with quantifier meanings. Let us assume that evaluative adverbs like *surprisingly*, *amazingly*, *fortunately* are sentence operators that predicate over a proposition.⁷ A sentence like *Surprisingly, a penguin is a bird* asserts that a penguin is a bird, while at the same time it conveys that (the speaker is of the opinion that) it is somehow unexpected that a penguin is a bird. Let us now look into what is expressed about expectation when such an adverb predicates over a quantified sentence. As observed above, (22a) suggests that fewer people were expected at the party, while in the case of (22b) more were expected.⁸

- (22) a. Surprisingly, a few people came to my party.
b. Surprisingly, few people came to my party.

In Nouwen (2005), I argue that evaluative predicates like *surprising* are monotone. That is, what is and isn't surprising is sensitive to entailments between propositions. If it surprises me to learn that Cindy came to Bill's party, then anything entailing that information should also surprise me. It cannot be that suddenly it does not surprise me that Cindy came to Bill's party late. In other words, *surprise* is downward monotonic: if *p* is

7. See Morzycki 2008; Katz 2005; Nouwen 2005 for a different use of such adverbs.

8. A reviewer notes that the point made below about (33) is really independent of the adverb *surprisingly*. S/he observes that even without such an evaluation, (33a) (resp. (33b)) would still indicate that more (resp. fewer) people than expected came. I agree that the adverbs are not crucial, but I disagree that the readings have to involve expectation. For instance, the same point as below can be made about what is desired instead of what is expected when *surprisingly* is replaced by *unfortunately*.

surprising, then anything entailing *p* is equally surprising.⁹ Applying this to (22) gives an interesting result. Upward monotonic quantifiers, as are downward monotone ones, are related by entailments. *A few people came to my party* is entailed by *Many people came to my party*, which in turn is entailed by *All people came to my party*. Similarly, *few people came to my party* entails that *not all people came* and is entailed by *no-one came*. This means that someone uttering (22a) is committed to being equally surprised had the number of people who came been many. Reversely, someone uttering (22b) is committed to being equally surprised had no-one come. To put it more succinctly, in (22a) the top end of the scale is marked as surprising, while in (22b) the bottom half is.

The above serves to illustrate that the formal properties of quantifiers, as provided by GQT, may help to explain why the negative/positive distinction manifests itself in the way it does. Does this mean that the denotational semantics of quantifiers may serve as a basis for a psychological model, where the monotonicity characteristics of a quantifier replaces the parameter of directivity? Despite clues that downward monotonicity and negativity are to some extent connected, the answer will have to be no. Monotonicity is too crude a measure to fully account for a much more complex phenomenon.¹⁰ First of all, directivity effects arise in interaction with context. Moxey (2006), for instance, shows that when a subject expects a quantity to be much larger than that expressed by a *positive* quantifier, this quantifier is sometimes taken to license complement anaphora. In her experiment, when a quantifier like *a small*

9. The claim that predicates like *surprise* are monotone is not easily defensible. This is because standard tests are useless due to several independent features of such predicates. (See Nouwen 2008b for details). Most importantly, *surprise* is factive. Consequently, (iiia) does not entail (iiib), for (iiia) does not entail that Billy wore a dress to the party. (An anonymous reviewer rightly pointed out that the role of factivity needs to be clarified. Thanks also to Crit Cremers and Øystein Nilsen for discussing this point with me on different occasions).

- (iii) a. I'm surprised that Billy came to the party.
- b. I'm surprised that Billy came to the party wearing a dress.

The monotonicity of *surprise* operates on a deeper level than what may become visible using tests like (iii). It is a constraint on what people's expectations/surprise may look like. If some actual or non-actual *p* (would) surprise(s) someone, then it cannot be the case that this surprise only holds for some particular worlds in which *p*. The surprise should hold for all normal situations in which *p* is true. Compare this with someone thinking it *possible* that *p*, for which such a constraint clearly does not hold. See Kadmon and Landman (1993) for another defence of the monotonicity of predicates like *surprise*.

10. This said, monotonicity has been successfully argued to account for various other aspects of quantifiers. For instance, Geurts (2003) shows that monotonicity properties of quantifiers play a role in syllogistic reasoning and Geurts and van der Slik (2005) show that monotonicity may explain some specific patterns of how quantified sentences are processed.

number of followed sentences expressing the expectation of *all*, the number of complement anaphora completions increased significantly. In other words, positive quantifiers can turn negative in certain contexts. Monotonicity is clearly insensitive to such considerations.

The shortcomings of monotonicity are further illustrated by comparing two specific monotone decreasing quantifiers. In Moxey et al. (2001), it is found that the monotone increasing quantifier *no less than 10%* gave, in a small number of cases, rise to negative perspective effects. In the same experiment, it was found that *at most 90%*, intuitively downward monotone, gave fewer negative-like responses than *no less than 10%* did.¹¹ Clearly, there is a variation in the data that is left unexplained if directivity is mapped onto the three-way classification of quantifiers into upward, downward and non-monotonic operators.¹² As we will now see, the generality and coarseness of the generalised quantifier approach has further drawbacks.

4.3 Further limitations of Generalised Quantifier Theory

Aside from applying formal tools like monotonicity to quantifier meaning, the ideas of generalised quantifier theory lead to a number of *big* questions, such as: Do all languages have these kinds of expressions? (see Barwise & Cooper 1981; Bach et al. 1995; Matthewson 2001) or: How many quantifiers can there be? (van Benthem 1986). Another important issue is the question of what distinguishes a possible quantifier, from an impossible one (Barwise & Cooper 1981; Keenan & Stavi 1986; van Benthem 1986). Barwise and Cooper (1981), for instance, claim that lexical quantifiers are always monotone (be it upward or downward). That is, non-monotone quantifiers are always derived, like *exactly twelve*, *between ten and thirty*, *more than six*, *but fewer than ten*, *precisely 50%*. Probably GQT's most well-known universal of this kind, however, is

11. Nouwen (2008c) argues that *no less than* quantifiers trigger implicatures which render them non-monotone. This, by itself, however, does not suffice to explain the observation, for the reasons explained in endnote 12.

12. A further complication is formed by non-monotone quantifiers. A sentence like *Exactly twenty percent of the students came to the party* licenses the inference that exactly eighty percent did not. Still, non-monotone quantifiers like *exactly twenty percent* are awful licensors of complement anaphora. This is consistent with data on upward entailing quantifiers that may give rise to a downward monotone implicature and thus end up being non-monotone. As an anonymous reviewer points out that *most* might give rise to a *not all* implicature, but this implicature does not suffice to license complement anaphora. Having said this, non-monotone quantifiers do behave as expected under evaluative predicates. An example like (iv) does not really fix the direction in which the amazement should be sought. For instance, (iv) could be because one expected more people, or fewer, or simply because it was Jasper's thirty-fifth birthday.

(iv) Amazingly, exactly thirty-five people came to Jasper's party.

the thesis that all natural language quantifiers (lexical and derived ones) are *conservative*. Conservativity is a property which precludes quantifiers from taking into account information about individuals that fall outside the domain of their first argument (the so-called domain of quantification). For instance, to know whether *every student is lazy*, it is completely irrelevant to look at individuals who are not a student. Investigating every student will do. The same goes for *most students are lazy*, *no students are lazy*, etc. If all your knowledge was restricted to properties of students, then you would still be able to decide on the truth-value of any sentence in which a quantifier is combined with the noun *student*.

There are, however, exceptions to the conservativity thesis, *only* being the most commonly discussed one. If you have just studied all the files of all the students of Utrecht University, and you know nothing else about the university, then you will not be able to decide on whether (23) is true or false. (The small caps indicate that *students* receives stress.)

(23) Only STUDENTS in Utrecht University are lazy.

The problem is that for (23) to be true, it is necessary that none of the non-students is lazy. So, you'd for instance need to study the professors as well in order to decide on (23).

Generalised quantifier theorists took this observation to mean that *only* is not a quantifier. This makes sense because of some other special properties *only* turns out to have. For instance, it is focus sensitive. If we change the stress in (23) from the noun to *Utrecht*, a totally different meaning emerges (which, for instance, entails that students in Leiden University are not lazy). Quantifiers like *every* are not sensitive to focus in this way. Another difference is that the distribution of *only* is much wider than that of *every* and its kin. *Only* combines with prepositional phrases (*only in AMERICA*), verb phrases (*he only LOOKS scary*), etc., while *every* is confined to its determiner position. It is clear then that *only* and *every* are simply different kinds of things: *only* is not a quantifier, *every* is. This begs the question, however, whether or not there are more things that look like a quantifier, but turn out to be something else. As it happens, the class of true natural language type (1,1) quantifiers is remarkably small.

Krifka (1999) was the first to point out that modified numerals are focus sensitive and that this is problematic for their treatment as type (1,1) quantifiers. For instance, with stress on the number word, a quantifier phrase like *at least four children* can answer a *how many* question. With stress on the noun, this is infelicitous.

(24) [Q] #Who left? / How many children left? [A] At least FOUR children left.

(25) [Q] Who left? / #How many children left? [A] At least four CHILDREN left.

If *at least four*, as a complex, is a relation between sets of individuals, then how is it going to be sensitive to focus on *four*? The least we should do is decompose this complex further.

Geurts and Nouwen (2007) observe that the focus sensitivity of modified numerals can lead to truth-conditional differences.

- (26) Cindy brought more than THREE bottles of wine. (She brought five.)
- (27) Cindy brought MORE than three bottles of wine. (She brought some beers too.)

Note that (26), but not (27) entails that Cindy brought (at least) four bottles of wine. In fact, (27) entails that what Cindy brought (properly) includes three bottles of wine. This is problematic: if *more than three* forms a quantifier, then *three bottles of wine* is not a constituent.

If *more than three* is not a type (1,1) quantifier, is *more than three N* then at least a type (1) quantifier? According to Hackl (2000), it is not. One of his arguments is based on examples like (28) and (29).

- (28) #Cindy separated more than one animal.
- (29) Cindy separated at least two animals.

To separate is compatible with objects whose cardinality exceeds one: somebody can separate two animals, but not one animal. Given this, the contrast between (28) and (29) is unexpected, since both *more than one animal*, and *at least two animals* are used for quantities of two animals or more. Instead of analysing *more than one* as a complex quantifier, Hackl's solution is to assume that *more than* is a proper comparative construction, with *more* being the comparative form of counting quantifier *many*. So, (28) is interpreted as [-er [Cindy separated d-many animals] [Cindy separated d-many animals & d = 1]]. The semantics of this says that the sentence is true if the maximal d such that Cindy separated d-many animals is larger than the maximal d such that Cindy separated d-many animals, where d is one. The example is unacceptable since one of the clauses in the comparison expresses that Cindy separated one animal. A similar structure for (29) is predicted to be felicitous, since in that case the comparison would contain a clause where Cindy separated two animals. Note that this solution is not available if we strive to view *more than one* as a type (1,1) quantifier.

The conclusion I want to draw from these issues is that it is a mistake to think that one can devise a general semantics for expressions of quantity. The reason this fails is that there is no coherent class of such expressions. Take modified numerals again. All modified numerals have in common that they contain a number word and some operator. It has become clear only very recently that there are very few other common characteristics shared by different kinds of modified numerals. For instance, Geurts and Nouwen (2007) argue that comparatively modified numerals (*more than n*) express a completely different kind of relation to a number than superlatively modified numerals (*at least n*) do. This might seem counterintuitive when examining simple examples like *Cindy ate more than three biscuits* or *Cindy ate at least four biscuits* (which will

seem equivalent), but a more detailed investigation uncovers sharp contrasts between the two kinds of quantifiers. The common factor in such contrasts, Geurts and Nouwen argue, is the intuition that *at least* expresses a modal attitude towards a quantity, while *more than* does not. This is illustrated in (30).¹³ While (30a) is a truism, (30b) is strange (or, some might say, false), since it suggests that there is the possibility of a triangle having more than three sides.

- (30) a. A triangle has more than two sides.
b. A triangle has at least three sides.

The negative counterparts of these quantifiers show similar effects: the word *bat* has fewer than four letters, but surely it is an inappropriate thing to say that it has at most three letters. The important point to note is that such data show that there are many ways in which a quantifier may communicate quantity information. Arguably, such specific ways are the consequence of the fact that many different kinds of constructions allow the expression of quantity. So, comparatively modified numerals are comparatives and so have the properties of a comparative; superlatively modified numerals are superlatives and will share properties with the class of superlative constructions, etc. This is made especially clear by prepositional numerals. Corver and Zwarts (2006) argue that quantifiers like *over a hundred* (as in *Cindy read over a hundred books last year*) is an adnominal PP. They moreover argue that prepositions in such quantifiers keep their normal spatial meaning, except that dynamic prepositions lose their motion sense. Nouwen (2008a), moreover, argues that some dynamic prepositions that lack a locative meaning have a fully directional meaning when they form a prepositional quantifier. The details of this argument are beyond the scope of the present paper, but, as an illustration, notice for instance that while *up to* places an upper bound on a quantity, it nevertheless is positive (in line with its upwards direction), this in contrast with other upper bound modified numerals.

- (31) Surprisingly, Cindy can be sentenced to up to six years in prison. less expected
(32) Surprisingly, Cindy can be sentenced to at most six years in prison. more expected

Putting it differently, variation in quantifier meanings is not simply due to a varied lexicon of quantifiers, but rather due to a variety of constructions that have the possibility of referring to quantities. The behaviour of the quantifier *up to six* should not be sought in a quantifier semantics, but rather in the normal spatial meaning of the preposition.

13. The modal nature of superlative modifiers is illustrated in many other ways too. See Geurts and Nouwen (2007) for the full details.

I started this section by saying that it makes sense to try and make precise what we mean with a (natural language) quantifier. It turns out that this is not a straightforward question to answer. There are probably only few true quantifiers in the type (1,1) sense (*every*, *no*, probably *most* and *some*).¹⁴ Attempts to generalise over a larger group of expressions of quantity miss out on distinctive semantic features. The key observation is that there is enormous variation in *how* quantities are expressed. Linguistic tools that exist independent of quantification may be exploited for that purpose (comparison, superlativity, spatial prepositions of direction, approximatives, etc.) The result sometimes resembles a type (1,1) quantifier, but in many cases, it does not.

5. Conclusion

The conclusion of the previous section was a rather negative one: the GQT notion of a quantifier is not really very suitable if we want to learn more about the semantics of expressions of quantity. If we want to appreciate the subtle differences with which quantifiers communicate quantities, a focus on how they differ is to be preferred over one which sets out to generalise as much as possible. While I showed that to some extent there is a relation between a notion like negativity and a formal property like downward monotonicity, such properties turn out too coarse to fit the data. So what's in a quantifier? Lots of stuff. The things we call *quantifier* are so varied, that they deserve to be studied on a case by case situation.

References

- Anscombe, J. & Ducrot, O. 1976. L'argumentation dans la langue. *Languages* 42: 5–27.
 Ariel, M. 2004. Most. *Language* 80(4): 658–706.
 Bach, E., Jelinek, E., Kratzer, A. & Partee, B. (eds). 1995. *Quantification in Natural Languages*. Dordrecht: Kluwer.
 Barker, C. 2002. The dynamics of vagueness. *Linguistics and Philosophy* 25: 1–36.
 Barwise, J. & Cooper, R. 1981. Generalized quantifiers and natural language. *Linguistics and Philosophy* 4(2): 159–219.
 Bradburn, N.M. & Miles, C. 1979. Vague quantifiers. *Public Opinion Quarterly* 43(1): 92–101.
 Corver, N. & Zwarts, J. 2006. Prepositional numerals. *Lingua* 116(6): 811–836.
 Dayal, V. 1998. *Any* as inherently modal. *Linguistics and Philosophy* 21: 433–476.
 Geurts, B. 2003. Reasoning with quantifiers. *Cognition* 86: 223–251.

14. One might think that GQT is thus reduced to a general theory of lexicalised quantifiers. This would lead to a very ad hoc characterisation of what a quantifier is, for it would have to exclude operators like *only* and *many/few* (cf. Solt 2006).

- Geurts, B. & Nouwen, R. 2007. At least et al.: the semantics of scalar modifiers. *Language* 83(3): 533–559.
- Geurts, B. & van der Slik, F. 2005. Ups and downs in syllogistic reasoning. *Journal of Semantics* 22: 99–117.
- Hackl, M. 2000. Comparative Quantifiers. Ph.D. dissertation, MIT.
- Hendriks, P. & de Hoop, H. 2001. Optimality theoretic semantics. *Linguistics and Philosophy* 24(1): 1–32.
- Horn, L. 2002. Assertoric inertia and NPI licensing. *CLS* 38(2): 55–82.
- Jayez, J. & Toven, L. 2007. Discourse inference and the meaning of *presque*. In *Proceeding of the ESSLLI 2007 Workshop on Quantifier Modification*, R. Nouwen & J. Dotlacil (eds). Dublin.
- Kadmon, N. & Landman, F. 1993. Any. *Linguistics and Philosophy* 16: 353–422.
- Katz, G. 2005. Attitudes toward degrees. In *Proceedings of SuB 9*, 1–3 November 2004, E. Maier, C. Bary & J. Huitink (eds). 183–196. Nijmegen: Radboud Universiteit Nijmegen.
- Keenan, E. & Stavi, J. 1986. A semantic characterization of natural language determiners. *Linguistics and Philosophy* 9: 253–326.
- Kibble, R. 1997. Complement anaphora and monotonicity. In *Proceedings of the Formal Grammar Conference*, G. Morrill, G.J. Kruijff & R. Oehrle (eds), 125–136. Aix-en-Provence.
- Krifka, M. 1999. At least some determiners aren't determiners. In *The Semantics/Pragmatics Interface from Different Points of View*, K. Turner (ed.), 257–291. Oxford: Elsevier.
- Kyburg, A. & Morreau, M. 2000. Fitting words: Vague language in context. *Linguistics and Philosophy* 23: 577–597.
- Lappin, S. 2000. An intensional parametric semantics for vague quantifiers. *Linguistics and Philosophy* 23: 599–620.
- Lewis, D. 1970. General semantics. *Synthese* 22: 18–67.
- Lindstroem, P. 1966. First order logic with generalized quantifiers. *Theoria* 32: 186–195.
- Matthewson, L. 2001. Quantification and the nature of cross-linguistic variation. *Natural Language Semantics* 9: 145–189.
- Montague, R. 1973. The proper treatment of quantification in ordinary English. In *Approaches to Natural Language*, K. Hintikka, J. Moravcsik & P. Suppes (eds), 221–242. Dordrecht: Reidel.
- Morzycki, M. 2008. Adverbial modification of adjectives: Evaluatives and a little beyond. In *Event Structures in Linguistic Form and Interpretation*, J. Dölling, T. Heyde-Zybatow & M. Schäfer (eds), 103–126. Berlin: Mouton de Gruyter.
- Mostowski, A. 1957. On a generalization of quantifiers. *Fundamenta Mathematicae* 44: 12–36.
- Moxey, L. 2006. Effects of what is expected on the focussing properties of quantifiers: A test of the presuppositiondenial account. *Journal of Memory and Language* 55: 422–439.
- Moxey, L. & Sanford, A. 1987. Quantifiers and focus. *Journal of Semantics* 5: 189–206.
- Moxey, L. & Sanford, A. 1993. Prior expectation and the interpretation of natural language quantifiers. *European Journal of Cognitive Psychology* 5: 73–91.
- Moxey, L. & Sanford, A. 2000. Communicating quantities: A review of psycholinguistic evidence of how expressions determine perspectives. *Applied Cognitive Psychology* 14: 237–255.
- Moxey, L., Sanford, A. & Dawydiak, E. 2001. Denials as controllers of negative quantifier focus. *Journal of Memory and Language* 44: 427–442.
- Nouwen, R. 2003. Complement anaphora and interpretation. *Journal of Semantics* 20(1): 73–113.
- Nouwen, R. 2005. Monotone amazement. In *Proceedings of the Fifteenth Amsterdam Colloquium*, P. Dekker & M. Franke (eds), 167–172. Amsterdam: ILLC.

- Nouwen, R. 2008a. Directional numeral quantifiers. In *Proceedings of SALT 18*. Ithaca NY: CLC Publications.
- Nouwen, R. 2008b. Degree modifiers and monotonicity. Paper presented at the workshop on Vagueness and Language Use, April 7–9, in Ecole Normale Supérieure, Paris.
- Nouwen, R. 2008c. Upperbounded no more: The implicatures of negative comparison. *Natural Language Semantics* 16(4): 271–295.
- O'Brien, B. 1989. Words or numbers? The evaluation of probability expressions in general practice. *Journal of the Royal College of General Practitioners* 39(320): 98–100.
- Pace, C.R. & Friedlander, J. 1982. The meaning of response categories: How often is 'occasionally', 'often', and 'very often'? *Research in Higher Education* 17(3): 267–281.
- Parikh, R. 1994. Vagueness and utility: The semantics of common nouns. *Linguistics and Philosophy* 17: 521–535.
- Paterson, K., Sanford, A. & Moxey, L. 1998. Quantifier polarity and referential focus during reading. *Journal of Memory and Language* 39: 290–306.
- Pracejus, J.W., Olsen, G.D. & Brown, N.R. 2004. On the prevalence and impact of vague quantifiers in the advertising of cause-related marketing (crm). *Journal of Advertising* 32(4): 19–28.
- Renooij, S. & Witteman, C. 1999. Talking probabilities: Communicating probabilistic information with words and numbers. *International Journal of Approximate Reasoning* 22: 169–194.
- Sanford, A., Dawydiak, E. & Moxey, L. 2007. A unified account of quantifier perspective effects in discourse. *Discourse Processes* 44(1): 1–32.
- Sanford, A., Fay, N., Stewart, A. & Moxey, L. 2002. Perspective in statements of quantity, with implications for consumer psychology. *Psychological Science* 13: 130–134.
- Sanford, A. & Moxey, L. 1993. *Communicating Quantities: A Psychological Perspective*. Hillsdale NJ: Laurence Erlbaum Associates.
- Sanford, A., Williams, C. & Fay, N. 2001. When being included is being excluded: A note on complement set focus and the inclusion relation. *Memory and Cognition* 29: 1096–1101.
- Schaeffer, N.C. 1991. Hardly ever or constantly? Group comparisons using vague quantifiers. *The Public Opinion Quarterly* 55(3): 395–423.
- Simpson, R. 1944. The specific meaning of certain terms indicating differing degrees of frequency. *Quarterly Journal of Speech* 49: 146–151.
- Solt, S. 2006. Why a few? and why not a many? In *Proceedings of Sinn und Bedeutung 10* [ZAS working papers in Linguistics 44(6)], C. Ebert & C. Endriss (eds.), 333–346. Berlin: ZAS.
- van Benthem, J. 1986. *Essays in Logical Semantics*. Dordrecht: Reidel.
- Wason, P. 1959. The processing of positive and negative information. *Quarterly Journal of Experimental Psychology* 11: 92–107.
- Westerståhl, D. 1984. Determiners and context sets. In *Generalized Quantifiers in Natural Language*, J. van Benthem & A. ter Meulen (eds), 45–71. Dordrecht: Foris.

Minimal versus not so minimal pronouns

Feature transmission, feature deletion, and the role of economy in the language system

Eric Reuland*
Utrecht University

In this article I start out from Kratzer's (2006) observation that a bound variable interpretation of German 1st and 2nd person (possessive) pronouns shows locality effects. I show how the approach based on syntactic chain formation developed in Reuland (2001, 2005a) resolves a general empirical problem in Kratzer's theory, while retaining the essence of her insight. I further show that the special interface condition on pronominal feature specification proposed by Kratzer is unnecessary. A general economy approach governing the division of labour between components of the language system is presented, which also subsumes Rule I as modified in Reinhart (2000, 2006).

1. Introduction

Over the last few years pronominal binding has given rise to a variety of intriguing puzzles. These puzzles range from the competition between variable binding and coreference as discussed in Reinhart (1983), Grodzinsky and Reinhart (1993), Reinhart (2000, 2006), to the competition between syntactic and semantic means for the encoding of anaphoric dependencies: feature checking/Agree in the sense of Chomsky (1995, 2005) versus variable binding in logical syntax as discussed in Reuland (2001, 2005a). Recently, Kratzer (2006) made an important contribution to this discussion. Part of her solution involves a very specific proposal about the competition between a morphosyntactic and a semantic process, which shares some insight with Reinhart's approach, but is less general.

In this contribution I will start out with a discussion of some of the particulars of Kratzer's proposal. I will show that the morphosyntactic part of her proposal runs into an empirical problem, but that the problem can be remedied in such a way that her

*I would like to thank the editors of this volume and two anonymous reviewers for their very helpful comments.

main insight is retained. Moreover, this modification makes it possible to dispense with her particular proposal about competition, bringing the facts she discusses directly under scope of economy views as developed by Reinhart and Reuland.

However, as noted in Reinhart (2006), the view of economy developed in Reinhart (1983), Grodzinsky and Reinhart (1993) and Reuland (2001) faces a number of problems. Consequently, Reinhart developed an alternative. In the final section of this article I will therefore discuss a number of basic issues around economy and arrive at a synthesis.

2. The starting point

One of the important issues in pronominal binding is how to account for bound variable (BV) reading of 1st and 2nd person pronominals, as exemplified in (1) (from Kratzer 2006):

- (1) I am the only one around here who can take care of my children.
- (2) Only you eat what you cook.

These sentences have both a strict and a sloppy reading. Limiting ourselves to (1), on the bound variable (sloppy) interpretation (1) implies that I can take care of my children and that that nobody else around here can take care of their own children. In the case of a strict reading, nobody but me can take care of my children. In a number of articles, Kratzer (1998, 2006) and Von Stechow (2003a, 2003b) address the puzzle of how to account for the sloppy reading. The source of the puzzle is that in the sloppy reading *my*, which is first person, must give rise to a variable that can range over individuals that are not. Kratzer and Von Stechow converge in assuming that the sloppy reading can only arise if the 1st person feature on the bound pronoun is not interpreted. They differ in the precise procedure. Kratzer assumes the existence of what she refers to as *minimal pronouns*, essentially pronouns with only a minimal set of morphosyntactic features, specifically, lacking a feature for person. Thus, the syntactic representation of (1) as it is interpreted is (3):

- (3) I am the only one around here who can take care of \emptyset children.

As Kratzer (2006) argues, in (1) the pronoun ends up being part of a syntactic agreement chain with its antecedent *I*. By this agreement it is supplied with a full feature specification from its antecedent, and is spelled out as *my* by a morphophonological spell-out rule. Von Stechow's analysis exploits the converse process of feature deletion. As part of the interpretation procedure at LF, features that cannot be interpreted are deleted under identity.

As both note, there is a high degree of intertranslatability between the two proposals. The main differences between the proposals reside in their assumptions about the

architecture of the linguistic system, and consequently, both in the motivation of the details of the proposals, and in the further predictions they make. In this contribution I will focus on the architectural issues, and show that the system developed in Reuland (2001, 2005a,b) captures the essential insights of both proposals, and resolves the problems that arise. I will start my discussion with issues around feature transmission and spell-out, and then come back to feature deletion.

3. Locality, agreement, and spell-out

Kratzer provides illuminating discussion of many issues that I cannot possibly go into here. I will limit myself to two issues. One is the contrast she discusses between her cases of (relatively) local binding and non-local binding, the other is the claim that features a pronoun acquires by agreement determine its realization. I will start with the former issue.

One of her important claims is that the dependency in (3) shows locality effects. This she shows on the basis of the following set of contrasts in German:

- (4) a. *Ich bin der einzige der meine kinder versorgt.
I am the only one who is taking care of my children.
- b. Wir sind die einzigen die unsere Kinder versorgen.
We are the only ones who take care of our children.

As I stated in the introduction, Kratzer's analysis is based on a process of feature transmission through an agreement chain. As she argues, transmission of features from *ich* to *meine* is blocked in (4a) due to the intervening agreement pattern between *der* and the explicit non 1st person *versorgt*. In (4b), due to syncretism *versorgen* is compatible with 1st person plural, hence transmission is not blocked.

Next, she makes another important observation: although feature transmission is subject to locality, not all cases of bound personal pronouns show locality. This is shown in (5):

- (5) a. *Du bist der einzige der deinen Vortrag versteht.
You are the only one who understands your talk.
- b. Du bist der einzige der glaubt dass jemand deinen Vortrag versteht.
You are the only one who thinks that somebody understands your talk.

Kratzer argues that long-distance cases such as (5b) involve logophoricity. She demonstrates a concomitant sensitivity to perspective as in (6):

- (6) a. I am the only one who wants DSS to take care of my children.
- b. *I am the only one who DSS wants to take care of my children.

In (6b) the perspective of the lower clause is DSS as the subject of the attitude verb *want*, which does not provide the value for *my*, hence the ill-formedness. To capture the contrast in (6) Kratzer introduces a semantics based on context shifting operators as proposed by Anand and Nevins (2004), which shift contexts to ones whose author or addressee is identified with the ‘origo’. Via the origo parameter, a believer’s individual alternatives can be plugged in directly as values for *de se* pronouns. Thus, in all languages a 1st person perspective operator appears in complements embedded by a 1st person attitude verb, and a 2nd person perspective operator in complements embedded by a 2nd person attitude verb.¹ As Kratzer states, the context shifting account produces apparent bound variable readings for 1st and 2nd person pronouns in languages such as German and English, even though they get their usual indexical interpretation. No syntactically represented operator-variable relation is needed at all.²

For details Kratzer refers to work in preparation, hence I will not pursue the implementation of the semantic strategy here, except for the fact that its existence is a crucial condition for Kratzer’s analysis to work.³

A core feature of Kratzer’s analysis is, then, that there are two strategies to derive sloppy readings: a syntax-based one, and a semantics-based one. As always, this situation raises the question of why the more general semantics based strategy cannot bypass a failure of the syntax-based strategy. The latter must assume that bound pronouns are generated with a minimal feature specification, and acquire the full specification needed for spell-out by agreement. The former must assume that it is possible to insert pronouns with a full specification right from the start. But if so, why cannot such a fully specified pronoun be inserted in (5a), bypassing the locality effects on agreement? Or, the way Kratzer puts it: If Minimal Pronouns can be born with certain features already in place, we expect there to be explicit guidelines regulating which features they have to have from the start, and which ones they can acquire via feature transmission. As she notes, this requires an interface strategy along the lines of Reinhart’s Rule I, and she proposes a version that performs the required role:

1. Amharic as discussed by Schlenker, and Zazakin as discussed by Anand and Nevis are special in that they also allow such operators embedded under a 3rd person attitude verb.

2. Since here the strategy is logophoric, no syntactic encoding is involved, and hence the explicit 3rd person feature of *who/wants* does not block this reading, only the fact that *who* – as a subject – carries the 1st person perspective matters. A reviewer notes that (6b) on the bound reading is not absolutely ungrammatical, but there there is indeed a clear preference for the non-bound interpretation, which is consistent with Kratzer’s point.

3. It is interesting to note, though, that Heim (2005) discusses a suggestion by Spector which does not involve transmission or deletion of features either. As Heim summarizes the proposal, all features manifest at PF (on DPs of semantic type *e*) are also present at LF and interpreted.

(7) Interface constraints for pronouns

- a. A pronoun is born with the minimal number of features that produces the intended interpretation.
- b. A pronoun acquires the maximal number of features that leads to a unique pronunciation.

It may be useful to give a brief recapitulation at this point of Reinhart's Rule I (Reinhart 1983). This rule regulates the availability of coreference versus variable binding in cases like (8):

(8) *Oscar admires him.

(8) is ruled out on the bound reading by condition B of the canonical binding theory. As we know bypassing the effect of condition B, using the option of assigning the pronoun *him* the same discourse referent as *Oscar*, should be blocked. This is what Rule I accomplishes:

(9) Rule I: Intrasentential Coreference

- NP A cannot corefer with NP B if replacing A with C, C a variable A-bound by B, yields an indistinguishable interpretation.

In (8) the coreferential and bound variable interpretation would be indistinguishable, hence coreference is blocked and condition B is free to apply.

Although in the spirit of Rule I, the formulation of (7) rather different. Since Kratzer presents no discussion of how (7) would capture the contrast between variable binding and co-reference Rule I was originally intended to capture – and it is indeed hard to see how it would – (7) appears to be specifically designed for the purpose of the contrasts discussed in this particular paper. Moreover, as Reinhart (2000, 2006) discusses, the original Rule I faced an important problem, as illustrated in (10):

(10) Max likes his mother and Felix does too.

This type of sentence allows both a strict reading, based on coreference, and a sloppy one, as indicated in (11a,b):

- (11) a. Max λx (x likes his mother & his = Max) & Felix λx (x likes his mother & his = Max)
- b. Max λx (x likes x's mother) & Felix λx (x likes x's mother)

The question is then why the strict reading is not blocked as it is in (8). One could say that the coreference option is licensed by the difference it makes for the interpretation in the second conjunct. Hence, making the interpretations distinguishable as in canonical cases such as *Obviously, everybody hates Oscar. Even Oscar hates him*. But, as Reinhart notes invoking this clause gives the wrong result for cases like (12):

(12) *Max praised him and Lucie did too (him \neq Max)

Only if *him* in (12) could be valued as Max, the second conjunct could mean that Lucie praised Max. So, by not having coreference in the first conjunct we influence the interpretive options in the second one, which is the same logic as invoked for (11), hence (12) should be good, but clearly it is not.

Because of this, Reinhart modifies Rule I, but very importantly for our present concerns, also reconceptualizes it. Reinhart originally stated the rationale behind Rule I in terms of economy: close an open expression as soon as possible. Since binding is a process closing a property, economy favours it to apply. Since this is the source of the problem, in Reinhart's new conception it is viewed as an effect of blocking: "If a certain interpretation is blocked by the computational system, you would not sneak in precisely the same interpretation for the given derivation, by using machinery available for the systems of use." (Reinhart 2000). This still rules out (8), but it does not enforce a bound variable interpretation whenever it is available as in (10). Kratzer's (7) does not take this reconceptualization into account. As I will show, upon proper consideration (7) can be dispensed with, and no more is needed than a generalization of Reinhart's Rule I.⁴

First, I will discuss the morphosyntactic part of the analysis: feature transmission and spell-out. Attractive, though this idea is, the proposal faces one major empirical problem. Many languages have anaphors that appear in the position where Kratzer would posit a minimal pronoun, and yet these are just spelled out as what they are: underspecified pronominal anaphors.

For instance, Russian (13) has the possessive anaphor *svoj*, which allows 1st, 2nd, or 3rd person antecedents (*ja*, *ty*, *on*, etc.) and exhibits the bound variable reading, but has the same form, irrespective of the phi-features its binder has – it is sensitive to phi-features of the head of the DP, though. This is in fact one more indication that a spell-out in terms of features of the binder is on the wrong track.

4. The discussion as I present it here contains a simplification. As is discussed extensively in Reinhart (2000, 2006) and references cited there, the strict/sloppy contrast also arises when the antecedent itself is a bound variable. This is illustrated by the two readings of (i):

- (i) Every wife thinks that only she respects her husband.
 - (a) Every wife (λx (x thinks that [$\text{only } x$ (λy (y respects x 's husband))]))
 - (b) Every wife (λx (x thinks that [$\text{only } x$ (λy (y respects y 's husband))]))

Consequently, what is involved is not co-reference in a strict sense, but rather the more general notion of co-valuation. The same issue comes up in (ii), given by an anonymous reviewer.

- (ii) Every boy thinks that he loves his dog and that St. Francis does too.

This type of facts bears on the details of the particular formulation Rule I receives in Reinhart's, system but not on the discussion of the rationale behind it. Consequently, it is not necessary to pursue this issue here.

- (13) Tol'ko ja ljublju/ty ljubish/on ljubit svoix detej.
only I love/you love/he loves POSS (BV) children

An object reflexive such *sebj**a* is also insensitive to phi-features of its binder. Dutch *zich*, German *sich*, etc. and Scandinavian possessive anaphors go against Kratzer's approach in the same way. Here too, there is no reflex of the phi-features of the antecedent and/or the corresponding pronominals on these anaphors when they are bound. This shows that in fact there is no general spell-out process of features acquired through transmission. Consequently, the relation between chain formation and spell-out must be reconsidered. If so, this challenges the idea that morphologically specified pronouns are inserted with minimal feature specification. But then the question comes up how the syntactic dependencies Kratzer showed to exist are implemented in the system, and why they show locality effects.

This might be taken to indicate that an approach based on LF deletion of features, like Von Stechow's might have a better chance of success, under the assumption that PF just spells out those features that are present in the numeration, and consequently, there is only one source in the lexicon for each pronominal. Hence, let us consider it.

4. A feature deletion approach

An approach based on feature deletion raises the question of how to restrict it. For instance, it should be prevented from applying in all the cases where Kratzer shows that locality conditions block feature transmission. Moreover, just like Kratzer's proposal needs a trade-off between a syntax-based and a semantics-based strategy, any restrictive implementation of Von Stechow's approach would require a similar Rule-I type strategy. Furthermore, a deletion analysis requires a principled answer to the question of what allows deletion. Even if a certain feature is not interpreted, prior to insertion it does not carry on its sleeve that deletion will be its fate. And clearly no grammatical theory can allow deletion of a potentially interpretable feature just for the sake of convergence.

I will now propose that the process of chain formation proposed in Reuland (2001) provides the basis for a synthesis.

Reuland (2001) addresses the question why many languages allow local binding of 1st and 2nd person pronominals while disallowing local binding of 3rd person pronominals. So, in Dutch we have paradigms such as the following:

- (14) a. Ik gedraag me goed.
I behave me well
b. Jij gedraagt je goed.
You behave you well
c. Hij gedraagt zich/*hem goed.
He behaves SE/him well


- (15) a. Ik voelde me wegglijden.
 b. Jij voelde je wegglijden.
 c. Hij voelde zich/*hem wegglijden. (*hem* allowed if disjoint)

It will be clear that semantically the pairs ⟨ik, me⟩, etc. in (14) must yield one argument. (For discussion of the difference between (14) and (15) I refer to Reuland 2001). This is effected by chain formation, which is based on the feature checking mechanism of Chomsky (1995).

Chomsky (1995) limits the means available to express dependencies within C_{HL} :

- (16) i. The only ways to establish a dependency: Move/Attract and checking.
 ii. The only way to force a dependency: checking grammatical features in a checking configuration.

The dependency between an anaphoric expression E and its antecedent is encoded by composing independently existing checking relations.



Just in case E has the appropriate feature composition (a SE-anaphor, or equivalent) composing R1, R2 and R3 may yield a composite dependency (DP, E), which carries over as an interpretive dependency at the relevant level. In the system of Chomsky (1995) checking configurations arise by feature movement, which is governed by attraction as in (22) (Chomsky 1995: 297):

- (18) K *attracts* F if F is the closest feature that can enter into a checking relation with a sublabel of K.

There is no separate checking operation. Being in a checking relation just has the following automatic effects:

- (19) a. a checked feature is deleted when possible
 b. a deleted feature is erased when possible

Possibility is relative to other principles of UG, specifically the principle of recoverability of deletion (PRD). Being in a checking configuration entails deletion/erasure up to recoverability. (19) expresses that checking takes place as soon as a checking configuration has been established. The consequence of being in a checking relation is that as many features are deleted/erased as possible. The basic claim in Reuland (2001) is (20):

- (20) The core syntactic dependency is that between two ‘occurrences’ of a feature, one being deleted/erased by the other.

This encodes a dependency since the content of a feature or feature bundle must always be recoverable when it is deleted. But this also entails (21):

- (21) Nothing bars deletion of an interpretable feature if no violation of the PRD results.

This is crucial for the understanding of 1st and 2nd person pronouns (and Nouns functioning as such). As shown in Reuland (2001), the structural Case relation between V and SE ensures that the formal features of SE are available on V, the relation between V and the T-system ensures that the features of SE, along with the features of V are visible in the T-system, and the latter enters into checking relation with the subject DP (Agreement), such that the subject DP when it checks the T-system for agreement features will also see the features of SE, check them and establish a dependency. Note that SE-anaphors are highly underspecified. They are only marked for person, or, more precisely, they are marked such that they are neither 1st nor 2nd person. As will be shown in more detail below, defining a dependency including a person feature does not violate the PRD. Thus, the procedure entails that SE-anaphors enter into a real dependency with their local antecedents within C_{HL} . The same holds true of 1st and 2nd person pronouns.

Chomsky (1995) restricts deletion to uninterpretable features, Reuland (2001) argues that no further condition on deletion is needed than the one that is independently given: deletion is only subject to the principle of recoverability of deletions (PRD). Deletion in a feature checking configuration should occur up to recoverability.

The intuition underlying the system is that a recovery relation establishes a dependency. Although at the time it was not put in that way, it quite similar to the intuition in Pesetsky and Torrego's (2004) system in which Agree/valuation turns different occurrences of a feature into different instances of that feature in a feature chain.⁵ And this is the way in which it is formulated in Reuland (2005a), which gives an implementation of binding chains in terms of Agree. As argued there, the properties of the head

5. The PRD has its roots in the earliest stages of generative grammar, when control was represented by "EQUI-NP deletion", deriving e.g. *He promised [PRO to come]* from *He promised [he to come]*. The rightmost occurrence of *he* was deleted under identity with the leftmost occurrence. Here the deletion is recoverable. The PRD expressed that *John is eating* should not be derived from *John is eating an apple*, with deletion of the object, since here the content of the object is not recoverable after the deletion. Note, that strictly speaking the PRD was not really satisfied in these examples, since two occurrences of *he* generally can, and in fact usually must receive different interpretations, one of the reasons why EQUI-NP deletion was rightly abandoned. So, what is required, is not just identity in form but also in possible contribution to interpretation in the given context. So, in order to meet the PRD the occurrence F^a_i and the occurrence F^a_j of feature F^a should not be able to make different contributions to interpretation in a given context; only then can they be unified as instances of this feature in a chain. As argued in more detail in Reuland (under contract) unification is available for different occurrences of D-features, person features, but not to number features which show some interesting interaction with gender that it would lead too far to discuss here.

(the C-T system and the v-V system), play a crucial mediating role in establishing the dependency. This is consistent with the facts that come up in Kratzer's discussion of the examples in (4). It is non-distinctness that allows chain formation to go through.⁶

The core idea is that unification of different occurrences into a chain, is allowed just in case PRD is not violated.⁷ In the case of for instance 1st person this is allowed, since all occurrences of a 1st person pronoun in one reportive context are interpreted identically. This is what is captured by Anand and Nevis's context operators as adopted by Kratzer. The same intuition is also captured by more syntactic approaches. For instance, Baker (2008) claims that 1st and 2nd person pronouns are bound by an operator in C. PRD prohibits deletion/unification where identity of interpretation is not guaranteed by inherent properties of a feature.⁸ In Reuland (2001) it was argued that grammatical number is a feature that resists deletion/unification. It was argued that grammatical number is found in 3rd person, whereas 1st and 2nd person exhibit inherent number. The latter notion is more precisely expressed in Kratzer's (2006) discussion of plurality and groups.⁹

6. Thanks to an anonymous reviewer for bringing this up.

7. An anonymous reviewer wonders what this implies for the analysis of SELF-anaphors, and for the status of the +R-feature in Reinhart and Reuland (1993). As already noted there, the $\pm R$ property of a pronoun is not a primitive, but to be explained in terms of independent properties its feature composition, and the interaction between these features and the linguistic environment. So, there is no sense in which a -R-feature unifies with a +R-feature. For SE-anaphors this explanation was achieved in Reuland (2001) and subsequent work. As outlined in Reinhart and Reuland (1993), Reuland (2001), and Reuland (2005b) SELF-anaphors are anaphoric due to the properties of SELF, which adjoins to and reflexive-marks the predicate. This leaves open a number of interesting issues for English reflexives, specifically why in the canonical environments a complex reflexive is required, even in 1st and 2nd person with intrinsically reflexive verbs. As argued extensively in Reuland (under contract), what is needed is one stipulation, namely that verbs like *behave* optionally project a syntactically transitive frame (with ACC case). In environments barring SELF movement, such as the subject position of ACC-ing, bound pronominals are actually widely allowed.

8. This formulation of course presupposes that features have properties. I believe they do. This entails that morphosyntactic features are linguistic objects, not just properties of linguistic objects. And, in turn this has non-trivial consequences for what a theory of features should look like. This issue is on my agenda, but discussion at this point would lead us too far a field.

9. For reasons that would lead me beyond the scope of the present contribution, I currently feel that it is worthwhile investigating whether a definiteness feature, as discussed by Kratzer, is involved in blocking local binding of 3rd person pronominals, perhaps in addition to the number feature, since there is interesting independent evidence from language change indicating that anyway number plays a role as well. See Postma (2004).

Given all this, let's come back to the difference between morphosyntactic and semantic encoding. There is intriguing evidence from Brazilian Portuguese (Menuzzi 1999) that a distinction between semantic binding and morphosyntactic syntactic encoding of a dependency is independently needed for personal pronouns. As Menuzzi shows, BP has two ways of expressing the 1st person plural pronominal, namely as *a gente* and as *nós*. 1st person interpretation notwithstanding, *a gente* is formally 3rd person, as indicated by verbal agreement. That is, *nós* and *a gente* differ in phi-feature composition.

As shown in (22), *nós* is a possible binder for *a gente* and vice versa. So, for binding a match in semantic type suffices.¹⁰

- (22) a. Nós achamos que o Paolo já viu a gente na TV.
We think that Paolo has already seen us on TV.
b. A gente acha que o Paolo já nos viu na TV.
We think that Paolo has already seen us on TV.

(23) shows that this option also exists in the more local environment of locative PPs.¹¹

- (23) a. Nós tínhamos visto uma cobra atrás de nós.
We had seen a snake behind us.
b. A gente tinha visto uma cobra atrás de nós.
We had seen a snake behind us.
c. A gente viu uma cobra atrás da gente.
We saw a snake behind us.

Dutch and English show no complementarity between pronominals and anaphors in this environment. That is, it is an environment outside the domain of the chain condition. But, interestingly, in the environment of (24) a semantic match is not sufficient. In this environment binding is ruled out unless antecedent and pronominal match in phi-features. *A gente* cannot bind *nos*, nor can *nos* bind the 3rd person clitic *se*, which would be the proper bindee for a *gente*.

- (24) a. Nós devíamos nos preparar para o pior.
We must prepare ourselves for the worst.
b. *A gente devia nos preparar para o pior.
c. A gente devia se preparar para o pior.
d. *Nós devíamos se preparar para o pior.

10. Just as in other languages a formal neuter – e.g. Dutch *het meisje* (*the girl*) – can bind a feminine pronominal, if there is a natural gender match.

11. Menuzzi does not give the *Nós .. a gente* pattern in PPs, but a Google search also instantiates this pattern.

This is precisely the domain of the chain condition. Since syntactic chains are based on phi-feature sharing, non-matching features result in an ill-formed syntactic object. The question is now why the semantic binding option, which (23) shows to exist, does not simply hide the unavailability of chain formation. Such hiding would be prevented if one were to assume that, extending Reinhart's original conception of economy, morphosyntactic encoding is preferred over encoding in the semantic system. However, given the problems with this view, an alternative is in order. Such an alternative is in fact provided in Chomsky (1995), see also Reuland (2001). Chomsky (1995) argues that ill-formed chains lead to a cancelled derivation. Within the economy approach of Chomsky (1995) a cancelled derivation blocks alternatives. That is, the organization of the grammar makes it impossible to bypass the chain-forming mechanisms, and take the option in which *a gente* in (18b) simply semantically binds *nos*.

This type of economy comparison falls squarely within Reinhart's revised rationale for Rule I: if your computational system rules it out, don't try another strategy to bypass that (see Section 5 for further discussion).

Consider again the possessive cases discussed by Kratzer. As she notes, these possessive pronouns are in a position where other languages allow possessive reflexives. Although they are not discussed in Reuland (2001), the same mechanisms of chain formation will be available to elements in this position, assuming that the outer layer of a DP is accessible for Agree/checking. The role of the phi-features of the inflected verb in mediating chain formation is consistent with the analysis of Reuland (2005a).¹² Along these lines, the approach developed for argument personal pronouns is available for possessive pronouns as well.

If so, the Chomsky-Reinhart rationale for comparing derivations will work as given, and prevent the semantic/logophoric strategy from bypassing what the ill-formedness of the ensuing chain would make impossible to encode in syntax.¹³

12. An issue yet to be resolved, though, is what factor determines whether a language has possessive anaphors or possessive pronouns. A quick survey suggests that languages marking definiteness by a pronominal article lack possessive anaphors (compare Dutch, English, German, Frisian, Modern Romance languages, etc. to Scandinavian (postnominal clitic), Latin (no marking), Slavic (no marking), etc. This suggests that it is the definiteness feature that blocks an anaphor in this position. This merits further investigation.

13. As always matters of execution remain to capture cross-linguistic variation. For instance, an anonymous reviewer raises the issue of how precisely chain formation is implemented in the case of an Icelandic Dative subject binding a singular possessor of a plural Nominative object. Also the singular-plural contrast in Danish possessives discussed in Safir (2004a) is interesting to study, especially in relation to the yet different behavior of Jutlandish dialects. Such issues, however, will have to be postponed to another occasion.

It is an interesting consequence of this line of thought that hard blocking effects are only to be expected where a morphosyntactic computation is in principle available, and consequently a categorical violation may ensue. In a language where the morphosyntactic computation is in principle not available, for instance due to peculiarities of its Case and agreement system, one will not expect to find categorical violations since there is no syntactic computation to be bypassed (as for instance in the case of local binding of object pronominals in Frisian, as discussed in Reuland and Reinhart 1995, see Reuland forthcoming for more extensive discussion). If so, one would expect no contrasts in the English counterparts of (4a) and (5a). As Kratzer discusses, there are contrasts in English, but weaker.¹⁴ This indicates that yet a different factor may be involved. This will be discussed in the next section.

As a final remark in this section, I wish to briefly address an issue that Kratzer raises at the end of her paper, namely the status of the bound variable interpretation of 3rd person pronominals. She raises the possibility that the absence of locality effects on the binding of 3rd person pronouns is the result of what I would like to refer to as masking: in addition to a locality sensitive real binding strategy for pronominals there is another strategy available treating and interpreting them as descriptive (D-type) pronouns.¹⁵ This is illustrated by the pair in (25):

- (25) a. None of these men was on vacation when the boss asked you where he was.
b. None of these men was on vacation when the boss asked you where the man was.

As she notes, if a definite description can produce an apparent bound variable reading, there is no reason why a D-type pronoun couldn't. Reuland and Avrutin (2004) contains an extensive discussion of the difference between what is called *pseudo-binding* and real variable binding on the basis of contrasts in dependencies involving preposed adverbial clauses. Although Reuland and Avrutin do not discuss locality contrasts, by itself their results show a clear difference between two types of dependencies that correlate with the distinction drawn by Kratzer. This shows that is very important to develop a proper conception of the role of economy in the language system. This is what I will do in the final sections of this contribution.

5. Rationales for economy

Reuland (2001) argued that the complementarity in Dutch and many other languages between locally bound SE-anaphors and related pronominals follows from economy.

14. See also Cable (2005), who argues for a different line, though.

15. Note, that masking requires that the derivations are not compared for economy.

I will briefly recapitulate the issue as it was formulated there. It is essentially the same problem as the one that came up in Kratzer's discussion. Consider the contrast in (26):

- (26) a. **Oscar* voelde *hem* wegglijden.
 Oscar felt him slide away.
 b. *Oscar* voelde *zich* wegglijden.

The puzzle it poses is that there is no reason to assume that (26a) directly involves a syntactic crash. Although the number and definiteness features of *hem* prohibit it from entering a chain with *Oscar*, it still has to be understood how this prevents just construing *hem* as a variable bound by *Oscar*. This would yield the same reading as (26b), given in (27):

- (27) *Oscar* λx (x felt [x slide away])

In Reuland (2001) this prohibition is derived by Rule BV, repeated here as (28). Rule BV is modelled on Reinhart's Rule I, and represent a general principle of economy of encoding.

- (28) Rule BV: Bound variable representation¹⁶
 NP A cannot be A-bound by NP B if replacing A with C, C an NP such that B heads an A-CHAIN tailed by C, yields an indistinguishable interface representation.

Standard rules of interpretation lead from (26b) to (27), but also from (26a) to (27). (28) expresses that *hem* in (26a) cannot be A-bound by *Oscar* due to the possibility for the alternative, *zich*, to enter a chain headed by *Oscar*.

Together with Rule I, Rule BV expresses the following hierarchy in the economy of encoding:

- (29) Economy of encoding:
 Narrow syntax < logical syntax (C-I interface) < discourse

Given the discussion of economy in Reinhart (2000, 2006) it is important to go carefully over the question of what aspects of mental computation could be reflected in such an economy metric.

16. For current purposes the simple version will do. Compare the full version: T may not translate an expression E' in SEM' with syntactically independent NP's A' and B' into an expression E in SEM in which A is A-bound by B, if there is an expression E'' resulting from replacing A' in E' with C', C' an NP such that B' heads a A-CHAIN tailed by C' and T also translates E'' into E.

(30) Possible rationales for economy

1. Demand on processing resources: close an open expression as soon as possible (Reinhart 1983)
2. Intrinsic (narrow syntax is automatic, blind)
3. Number of crossmodal steps (Reuland 2001)
4. Temporal course of parsing (Friederici & Kotz 2003)
5. Cooperation: "If a certain interpretation is blocked by the computational system, you would not sneak in precisely the same interpretation for the given derivation, by using machinery available for the systems of use." (Reinhart 2000)

I will now discuss these rationales in turn:

1. Reinhart (1983) proposed this rationale in her discussion of Rule I. It appears to reflect a preference of the language system for variable binding over co-reference. In a conception where pronouns are construed as variables, expressions with free pronouns are essentially open expressions that become closed only after pronoun is valued. Variable binding entails closing an open expression immediately, whereas coreference allows closing the expression only after a further search through the discourse store. Under the assumption that keeping an expression open requires keeping it in working memory, and that keeping material in working memory carries a cost, variable binding is less costly, hence preferred.
2. Another possible rationale is that certain subsystems of the language system are intrinsically cheaper in their demand on processing resources than others. For instance, in the case of narrow syntax one could argue that many, if not all of its operations are forced by triggering features. In operations such as Movement, Case checking, or Agreement there is no choice. In the interpretation of *himself* in *John was standing below himself* we have no alternative but concocting an impossible situation in order to interpret it, the syntactic machinery is fully deterministic and leaves no options. Automatic, blind processes require little monitoring while they apply, you only have to worry about the outcome. This certainly is a possible rationale for why narrow syntax would be economical. Finding a binder for a variable is less constrained in logical syntax, hence more monitoring and attention is required, which would make it more costly. The same holds true *a fortiori* for the discourse system, which therefore would come out as costliest.
3. The discussion in Reuland (2001) takes human laziness as a rationale. If you are doing something, the easiest thing is to just keep doing it. This translates as a cost associated with cross-modal steps, which require starting to do something else. So, going from two basic expressions to a coreferential interpretation requires 4 cross-modal steps, from to a BV interpretation 3, and interpreting them via syntactic chain formation requires only two cross-modal steps. The result is an economy metric that can be applied without taking grammar-external operations into consideration.

(31) a.	Discourse storage (values)	a		a
	C-I objects (variables)	x ₁		x ₂
	Syntactic objects (CHAINS)	C ₁		C ₂
	Basic expressions	α	...	β
b.	Discourse storage (values)	a		
	C-I objects (variables)	x ₁	←	x
	Syntactic objects (CHAINS)	C ₁		C ₂
	Basic expressions	α	...	β
c.	Discourse storage (values)	a		
	C-I objects (variables)	x ₁		
	Syntactic objects (CHAINS)	C ₁	←	C ₂
	Basic expressions	α	...	β

4. Work on the time course of natural language parsing by Friederici and Kotz (2003) provides yet a different rationale: irrespective of intrinsic cost of steps or components of steps. It is obvious that in order to be able to even look for the value of an expression in discourse, the expression must be there in at least some rudimentary form. Variable binding relations require the representation of at least some basic syntactic hierarchy. Otherwise they are impossible to compute. Assigning a discourse value requires that at least the nature of the objects that would qualify is determined. Thus, the logic of interpretation imposes an intrinsic order on the type of processes that must be applied.¹⁷

Although there is considerable debate on the issue of the time course in the processing literature, not all of it is based on a sufficient analysis of the logical structure of the task or on tasks with sufficient temporal resolution. That is, in order to claim that discourse information is used simultaneously with grammatical information one should have tasks that that sufficiently discriminating in temporal resolution, and should address the issue how discourse information can be coupled with chunks of

17. They present a time course model that relates differences in the stages of the interpretation procedure to differences in ERP signature and in some cases to differences in the fMRI signal. Discussion here would lead us beyond the scope of this work. The interested reader is referred to the works cited.

the signal whose category and constituency have not yet been determined. Koornneef (2008) presents an extensive discussion of the processing literature and he shows that some results in the literature (Badecker & Straub 2002), that appear to be inconsistent with a temporal ordering as sketched, are arguably lacking in temporal resolution.

5. As already discussed in Section 2, Reinhart (2000, 2006) shows that one must be careful to avoid a too simplistic interpretation of economy. From the pattern represented in (10), (11) and (12) Reinhart concludes that something is wrong with the economy based interpretation of Rule I. As she states it in Reinhart (2006: 183): “However, as plausible as the “least effort” approach seems, it is not clear to me that the human processor is indeed sensitive to this type of economy considerations.” Instead she proposes the rationale quoted above, repeated here as (32):

(32) *Against sneaking in*

“If a certain interpretation is blocked by the computational system, you would not sneak in precisely the same interpretation for the given derivation, by using machinery available for the systems of use” (Reinhart 2000).

In Reinhart (2006: 185) she argues that this rationale follows from a broader economy principle ‘minimize interpretive options’, a principle that prohibits applying a procedure that increases the number of interpretations associated with a given PF.

However, in assessing the issue we should distinguish between preference and blocking. By now there is a considerable amount of converging evidence from processing experiments of various kinds that a bound variable interpretation is more easily accessed than a coreferential interpretation. Burkhardt (2005) presents a range of processing experiments on English to this effect. Vasic (2006) reports systematic differences in the performance on strict versus sloppy readings by agrammatic aphasics, showing that they have considerable trouble with strict readings. Koornneef et al. (2006) reports eye-tracking experiments showing that even in contexts pragmatically biased for a strict reading the initial reading is the bound variable one, followed by a repair action. As a consequence, one must conclude that the human processor is indeed sensitive to some sort of economy consideration as discussed, even if one leaves open the precise source (as in 1, 2, 3, or 4), see Ktoornneef (2008) for more experiments and discussion. Nevertheless, the problem noted by Reinhart for the VP-ellipsis cases is real and should be resolved. This entails that both types of considerations are correct, and that both are reflected independently in Kratzer’s data: the categorical contrast in German and the weaker contrast in English. I will now show that a solution of the problem is to be found in combining the two approaches: ‘ranking procedures’ and the ‘prohibition against sneaking in.’

Note that even Reinhart’s rationale does not escape from minimally making reference to the time course of interpretation. The process of determining whether a procedure

increases the number of interpretations is defined with respect to a given representation as a starting point. (But of course, it is also compatible with stronger positions as in 2 or 3). Upon proper consideration there is another possible rationale for the prohibition to sneak in, which stays closer to the original intuition:

- (33) The grammatical system does not consider alternatives for impossible derivations: Rejection by one component of the grammar is final.

From a processing perspective (33) expresses the claim that entering an impossible derivation is equivalent to entering a garden path from which recovery is impossible. From a grammatical perspective, (33) subsumes the effects of a cancelled derivation in the economy ranking of Chomsky's (1995) system. If a set of alternatives contains a cancelled derivation - a derivation in which an impossible chain (with non-matching features) has arisen - the cancelled derivation still outcompetes the alternatives, which are therefore not being considered.

This is precisely what is needed to capture the facts from Brazilian Portuguese discussed in Section 2. But also note, that these facts support (33) as a rationale for the prohibition against 'sneaking in' over a purely cooperative strategy of minimizing interpretive options, since in fact the interpretive options are not really minimized. Consider again the crucial part of the paradigm repeated as (34):

- (34) a. Nós devíamos nos preparar para o pior
We must prepare ourselves for the worst
b. *A gente devia nos preparar para o pior
c. A gente devia se preparar para o pior
d. *Nós devíamos se preparar para o pior

A *gente* cannot bind *nos*, nor can *nos* bind the 3rd person clitic *se*, which would be the proper bindee for a *gente*. That is, in principle they could, but in the local domain a cancelled derivation results which block the binding alternative. But note, that in this case it would be ironic to speak of a strategy of minimizing interpretive options, since the result is no options. A better characterization is: You are not allowed to bypass a narrow syntax prohibition by resorting to a 'logical syntax' strategy. This is precisely what (33) expresses.

The Brazilian Portuguese facts show two important things: (1) There is evidence for a competition between narrow syntax and logical syntax, along the lines proposed in Reuland (2001). (2) A *rejection is final rationale is to be preferred over a cooperation rationale*.

The question is now how the prohibition in (33) tallies with the economy hierarchy of alternatives. In this Section I will argue that an economy hierarchy of alternatives and the prohibition in (33) can indeed be combined, and in fact jointly enable us to return to a quite conservative position with respect to economy of derivation.

6. Economy: A synthesis

As in Reuland (2001), I assume that three main components of the language system are involved, and claim that they ranked in an order that reflects intrinsic asymmetries in the way they are accessed in interpretation:

- (35) Hierarchy of access
 narrow syntax < logical syntax < discourse

That is, given some expression that is to be interpreted, minimal properties of syntactic structure must have been established before semantic processes can start operating, and minimal properties of semantic structure must have been established before any link can be established with discourse structure. Note, that this does not presume that the operation of one component must have been completed before another can start, only that it must have started before the next one can start. It does not presume either that the processor cannot entertain hypotheses about discourse, before either syntactic or semantic computation starts. We constantly develop expectations as to what is going to happen, whether you see clouds gathering, or a drunken driver approaching you. The claim is only that in order for such expectations to be able to bear on the course of interpretation, the structure that is being built must contain objects for these expectations to be connected to. Thus, the intrinsic order in which components of the language system become applicable determines a hierarchy of accessibility. For present purposes it is immaterial whether this order of applicability is based on only on time course, and intrinsic conditions on cross-modular steps, or also on intrinsic differences in the cost or processes, handling of more complex information, or all of these.

Yet, what matters is not just the intrinsic order. Also economy comes into play. In one of its roles it has much of its original flavour:

- (36) Relative Economy
 As soon as a component becomes applicable, economy favours its processes to complete their operation.

If not, the system could always wait for more information to become available. As for instance the eye-tracking results obtained by Koornneef (2008) show, this is not what happens. For instance in the case of semantics versus discourse, once the syntactic structure has sufficiently developed for semantic processes to start operating, the processor does not wait until it is sure about what to do. Rather, it first creates a variable binding reading, and only then reverts to the contextually favoured coreference reading. This notion of economy can be grounded in the very considerations proposed in Reinhart (1983): It is costly to maintain an open expression. It is this relative economy one may assume is operative in the case of the preferred interpretation English bound variable pronouns as discussed in Section 3 and 4.

The intrinsic order of access is coupled with the generalized principle restricting accessibility we already discussed:

- (37) Absolute economy: Rejection is final

If the derivation of a particular interpretation of a certain expression is blocked in a given component of the language system this derivation is cancelled (hence access to subsequent components in the hierarchy to derive precisely the same interpretation for the given expression is rendered impossible).

Consider again the *zich/hem* contrast in (26) and (27), repeated here as (38) and (39):

- (38) a. **Oscar* voelde *hem* wegglijden
Oscar felt him slide away
b. *Oscar* voelde *zich* wegglijden

- (39) *Oscar* λx (x felt [x slide away])

Given (37), the competition between *zich* and *hem* is no longer taken to be direct, as in Reuland (2001). It is indirect, and based on the following steps:

- (40) i. forming an $\langle \textit{Oscar}, \textit{zich} \rangle$ chain is allowed, given the principles of chain formation given
ii. forming an $\langle \textit{Oscar}, \textit{hem} \rangle$ chain is not allowed given the PRD
iii. directly forming (39) from (38a) in logical syntax is blocked by (37)
iv. only $\langle \textit{Oscar}, \textit{zich} \rangle$ is left as a source for (39)

In this derivation, the $\langle \textit{Oscar}, \textit{hem} \rangle$ chain is independently ruled out, not just under comparison with the $\langle \textit{Oscar}, \textit{zich} \rangle$ chain. Although on a meta-level we can still say that one derivation is more economical than the other, this economy is not measured in a direct comparison between derivations. The analysis is now consistent with a conservative notion of economy of derivation. No derivations need be compared with numerations containing different elements. It is not necessary either to limit numerations to lexical elements, as in Boeckx et al. (2007).¹⁸

18. Note, that (37) only comes into play when one derivation of a certain interpretation is blocked and an alternative strategy to bypass it should be prevented from applying. So, it says nothing about the two interpretations of Heim's (i), brought up by an anonymous reviewer:

- (i) Only I think that I am smart

No grammatical principle is violated in either of these derivations, and moreover the interpretations are indeed different. It does come up in (4a), though. Since, here, given Kratzer's assumptions, chain formation is an option in the general configuration, that fact that it is blocked by grammatical conditions on consistence of feature chains (37) prevents the same interpretation being assigned by different means. That is, it prevents the strategy

The present approach has another important consequence. In a language system allowing locally bound (3rd person) pronominals – as in Frisian – the absence of a competitor is not enough. If the pronominal could enter the chain – that is, if all conditions for chain formation were satisfied – and the resulting object would violate the PRD, (37) would simply block any alternative using the same element. In the absence of a competitor this would entail ineffability.¹⁹ Given this, what we expect is that locally bound pronominals are never an isolated phenomenon in a language. Their presence always involves more than just the absence of a competitor.

Frisian presents a crucial case to illustrate this point. As discussed in detail in Reuland and Reinhart (1995), based on independent evidence from Hoekstra (1994), a locally bound pronominal is licensed not via structural Case, but via inherent Case. This entails that the issue of forming a chain link does not arise; chain formation is not blocked due to a violation the ensuing chain would cause, but the conditions for chain formation are not met. Hence (37) does not come into play. Van Gelderen (2000) presents an extensive discussion of earlier stages of English, which also allows locally bound pronominals. She shows convincingly (see specifically Chapter 5) that the loss of this feature coincides with the independently established loss of inherent case.

It is a recurrent question whether an approach to variable binding along the lines of (37) is sufficient to capture more complex patterns of variable binding. In recent literature we find a debate about a puzzle that has been first noted in Dahl (1974). (See Roelofsen 2007 for a very useful discussion.) The question is how to derive the correct interpretive options in (41), that is what the possible values of *he* and *his* are in the italicized completion of the elided phrase:

- (41) Max said that he called his mother. Bob did too (= *Bob said that he called his mother*)
- a. . . . Bob too said that Bob called Bob's mother.
 - b. . . . Bob too said that Max called Max's mother.
 - c. . . . Bob too said that Bob called Max's mother.
 - d. . . . Bob too said that Max called Bob's mother.

that apparently is available in cases such as (5b) to apply in (4a), yielding the interpretation or interpretations that syntax blocks. Note, that, unlike what an anonymous reviewer suggests, no comparison between the coreference interpretation and the bound variable interpretation is involved, and no equivalence between the two is intimated.

19. As we will see, in some cases, local ineffability is indeed what a language system can deal with. If so, as an alternative a very different mode of expression may be used.

As Roelofsen puts it, the challenge is to account for the fact that (41-a), (41-b), and (41-c) are possible readings of the elided phrase, while (41-d) is not. The fact has been discussed by a variety of authors, including Fiengo and May (1994), Heim (1998), Fox (1998), Safir (2004b), and Reinhart (2006). A recurrent property of accounts proposed for this pattern is that somehow local binding of *his* by *he* must be enforced. To this end Fox (1998) formulates Rule H:

(42) Rule H

A variable x cannot be bound by an antecedent a , if a more local antecedent b could bind x yielding an indistinguishable interpretation.

In order to obtain the illegitimate (41d), *Max* must have been able to bind *his* and corefer with *he* in the antecedent clause, as in (43):

(43) $\text{Max } \lambda x (x \text{ said that he} = \text{Max called } x\text{'s mother})$

Rule H forbids that. With varying technical details also Safir (2004b) and Buring (2005) explore the intuition that there is a locality restriction on binding. Reinhart (2006) argues that no special condition is needed in the system she developed, but Roelofsen shows that in the specific elaboration she presents, in fact not only is (41d) correctly blocked, but also, incorrectly, (41c). On the basis of this, Roelofsen argues that one should apply Reinhart's insight in a simpler and more intuitive way. He proposes the following formulation:

(44) Rule S: a Simpler Interface Rule

Interpretations which are ruled out by restrictions on binding cannot be sneaked in via other anaphoric mechanisms.

It is easily seen that (44) is just an instantiation of (37). Under this formulation neither (41d), nor (41c) are blocked. Roelofsen then argues for a separate locality condition to rule out (41d). The question is, then, whether one should keep pursuing a unified solution for VP-ellipsis, including Dahl's puzzle, or be content with a two-component solution. Initially, one might be inclined to reject a two-component solution, but the issue is really empirical. Namely, do all the restrictions one has to account for indeed have the same status? In fact, I believe there is evidence that they don't. Canonical binding theory violations are categorical, and hard to impossible to obviate. However, the local binding requirement expressed by Rule H and its alternatives is quite sensitive to plausibility considerations. This is shown in the following example:

(45) Suppose John is a gambler, Bill does not like gambling himself but entrusts John with his capital.

- a. Dutch: Jan durfde niet toe te geven dat hij in Vegas zijn fortuin verspeeld had en Bill ook niet.

- b. English: John did not dare to admit that in Vegas he lost his fortune, and Bill did not either.

In both English and Dutch the conjunct is fine with a J J J B J B interpretation. Since we have to construct an appropriate context to obtain this interpretation, and it is hard to impossible to get it in (41), it seems safe to say that indeed some principle favouring local binding is active in language. However, its status must be that of a processing preference that is not grammatically encoded, and comparable to the preference for a sloppy interpretation over a strict interpretation in VP-ellipsis that has been demonstrated by psycholinguistic experimentation. My conclusion is then, that the interplay between economy considerations and binding possibilities is best captured by a hierarchy of subsystems as expressed in (35), with their workings governed by (36) and (37).

7. By way of conclusion

The following conclusions appear warranted:

- i. A chain formation approach along the lines of Reuland (2001, 2005) resolves a general empirical problem in Kratzer's transmission and spell-out theory, that is caused by under-specified anaphors, while retaining her essential insight that certain bound variable dependencies show locality effects.
- ii. The fact that the English counterparts of certain ill-formed German sentences are better than expected may be due to the fact the German encoding strategy is not available in English, hence does not lead to a Rule I effect.
- iii. There is no need to supplement Reinhart's (2000, 2006) rule I by a further interface condition on pronominal feature specification.
- iv. Rule I as modified in Reinhart (2000, 2006) can be subsumed under a general economy approach governing the division of labour between components of the language system.

References

- Anand, P. & Nevins, A. 2004. Shifty operators in changing contexts. In *Proceedings of SALT XIV*, R.B. Young (ed.). Ithaca NY: CLC Publications.
- Badecker, W. & Straub, K. 2002. The processing role of structural constraints on the interpretation of pronouns and anaphors. *Journal of Experimental Psychology: Learning, Memory and Cognition* 28: 748–769.
- Baker, M. 2008. *The Syntax of Agreement and Concord*. Cambridge: CUP.

- Boeckx, C., Hornstein, N. & Nunes, J. 2007. Overt copies in reflexive and control structures: A movement analysis. In *University of Maryland Working Papers in Linguistics* 15, A. Conroy, C. Jing, C. Nakao & E. Takahashi (eds), 1–46. College Park MD: UMWPI.
- Büring, D. 2005. Bound to bind. *Linguistic Inquiry* 36(2): 259–274.
- Burkhardt, P. 2005. *The Syntax–Discourse Interface: Representing and Interpreting Dependency*. Amsterdam: John Benjamins.
- Cable, S. 2005. Binding local pronouns without semantically empty features. Ms, Dept. of Linguistics, MIT.
- Chomsky, N. 1995. *The Minimalist Program*. Cambridge MA: The MIT Press.
- Chomsky, N. 2005. On phases. Ms, MIT.
- Dahl, Ö. 1974. How to open a sentence: Abstraction in natural language. In *Logical Grammar Reports* 12. Gothenburg: University of Gothenburg.
- Fiengo, R. & May, R. 1994. *Indices and Identity*. Cambridge MA: The MIT Press.
- Fox, D. 1998. Locality in variable binding. In *Is the Best Good Enough?*, P. Barbosa, D. Fox, P. Hagstrom, M. McGinnis & D. Pesetsky (eds). Cambridge MA: The MIT Press.
- Friederici, A.D. & Kotz, S.A. 2003. The brain basis of syntactic processes: Functional imaging and lesion studies. *NeuroImage* 20: 8–17.
- Gelderen, E. van. 2000. *A History of English Reflexive Pronouns* [Linguistik Aktuell/Linguistics Today 39]. Amsterdam: John Benjamins.
- Grodzinsky, Y. & Reinhart, T. 1993. The innateness of binding and coreference. *Linguistic Inquiry* 24: 69–101.
- Heim, I. 1998. Anaphora and semantic interpretation: A reinterpretation of Reinhart's approach. In *The Interpretive Tract: MIT working papers in Linguistics*, U. Sauerland & O. Percus (eds), 205–246. Cambridge MA: The MIT Press. (Written and distributed as a technical report at the University of Tübingen in 1992).
- Heim, I. 2005. Features on bound pronouns: Semantics or syntax? Ms, Department of Linguistics, MIT.
- Hoekstra, J. 1994. Pronouns and case: On the distribution of Frisian *harren* and *se* 'them'. *Leuvense Bijdragen* 83: 47–65.
- Koornneef, A.W. 2008. *Eye-catching Anaphora*. Utrecht: LOT Dissertation Series.
- Koornneef, A.W., Wijnen, F. & Reuland, E. 2006. Towards a modular approach to anaphoric processing. In *Ambiguity in anaphora workshop proceedings*, R. Artstein & M. Poesio (eds), 65–72. Malaga.
- Kratzer, A. 1998. More structural analogies between pronouns and tenses. In *Proceedings of SALT VIII*, D. Strolovitch & A. Lawson (eds), 92–110. Ithaca NY: CLC Publications.
- Kratzer, A. 2006. Minimal pronouns. Ms, Department of Linguistics, University of Massachusetts, Amherst.
- Menuzzi, S. 1999. *Binding Theory and Pronominal Anaphora in Brazilian Portuguese* [HIL dissertation: Leyden, LOT International Series 30]. The Hague: HAG.
- Pesetsky, D. & Torrego, E. 2004. The syntax of valuation and the interpretability of features. Ms, MIT & University of Massachusetts, Boston.
- Postma, G. 2004. Structurele tendensen in de opkomst van het reflexief pronomen 'zich' in het 15^{de} eeuwse Drenthe en de theorie van reflexiviteit. *Nederlandse Taalkunde* 9(2): 144–168.
- Reinhart, T. 1983. *Anaphora and Semantic Interpretation*. London: Croom Helm.
- Reinhart, T. 2000. Strategies of anaphora resolution. In *Interface Strategies*, H. Bennis, M. Everaert & E. Reuland (eds), 295–325. Amsterdam: Royal Academy of Sciences.
- Reinhart, T. 2006. *Interface Strategies: Optimal and Costly Computations*. Cambridge MA: The MIT Press.
- Reinhart, T. & Reuland, E. 1993. Reflexivity. *Linguistic Inquiry* 24: 657–720.

- Reuland, E. 2001. Primitives of binding. *Linguistic Inquiry* 32(2): 439–492.
- Reuland, E. 2005a. Agreeing to bind. In *Organizing Grammar: Linguistic Studies in Honor of Henk van Riemsdijk*, H. Broekhuis & N. Corver (eds). Berlin: Mouton de Gruyter.
- Reuland, E. 2005b. Binding conditions: How are they derived? In *Proceedings of the HPSG05 Conference Department of Informatics, University of Lisbon*, S. Müller (ed). Stanford CA: CSLI Publications.
- Reuland, E. Forthcoming. *Anaphora and Language Design*. Cambridge MA: The MIT Press.
- Reuland, E. & Avrutin, S. 2004. Binding and beyond: Issues in backward anaphora. In *Anaphora Processing: Linguistic, Cognitive and Computational Modelling*, A. Branco (ed). Amsterdam: John Benjamins.
- Reuland, E. & Reinhart, T. 1995. Pronouns, anaphors and case. In *Studies in Comparative Germanic Syntax*, H. Haider, S. Olsen & S. Vikner (eds), 214–269. Dordrecht: Kluwer.
- Roelofsen, F. 2007. Constraints on the interpretation of bound and referential pronouns. Ms, Institute for Logic, Language, and Computation, University of Amsterdam.
- Safir, K. 2004a. *The Syntax of Anaphora*. Oxford: OUP.
- Safir, K. 2004b. *The Syntax of (In)dependence*. Cambridge MA: The MIT Press.
- Stechow, A. von. 2003a. Feature deletion under semantic binding, S. Kawahara & M. Kadowaki (eds.), *Proceedings of the Thirty-Third Annual Meeting of the North East Linguistic Society*, 379–404.
- Stechow, A. von. 2003b. Postscript to feature deletion under semantic binding. <<http://www2.sfs.uni-tuebingen.de/~arnim10/Aufsaeetze/index.html>>.
- Vasić, N. 2006. *Pronoun Comprehension in Grammatic Aphasia: The Structure and Use of Linguistic Knowledge*. Utrecht: LOT International Dissertation Series.

Against partitioned readings of reciprocals

Sivan Sabato & Yoad Winter

The Hebrew University of Jerusalem/Utrecht University

This paper examines interpretations of sentences with reciprocal expressions like *each other* or *one another*. We concentrate on cases where two or more separate groups can be discerned in the interpretation of the subject of predication. We study the availability of such partitioned interpretations with definite subjects and proper name conjunctions, and show new evidence that partitioning effects are independent of the semantics of the reciprocal expression, and are exclusively determined by the interpretation of the subject. We then propose that the effect is yet another result of the familiar dependency of descriptions on contextual quantifiers.

1. Introduction: Partitioned interpretations of reciprocal sentences

Sentences with *reciprocal expressions* like *each other*, *one another* or *mutually* involve a variety of interpretations, which have been in the focus of much recent research. In this paper we examine a special kind of interpretations of reciprocal sentences, which we call *partitioned interpretations*. These interpretations involve reciprocal relations between two or more disjoint sets, with no reciprocal relations between the different sets. For example, sentence (1), from Fiengo and Lasnik (1973), has a partitioned interpretation because it is acceptable in the situation depicted in figure 1, where reciprocal hitting relations appear within two disjoint sets of men, but not between those sets.¹

(1) The men are hitting each other.

Another example of a partitioned interpretation is exemplified by sentence (2) below, from Dalrymple et al. (1998). This sentence is judged to be true when there are several disjoint stacks of planks. In such cases each stack is connected via the relation *be stacked atop*.

1. We are grateful to Tali Ore for creating this figure.




Figure 1. The men are hitting each other

- (2) He and scores of other inmates slept on foot-wide planks stacked atop each-other.

What are the origins of partitioning effects with reciprocals? A priori, partitioning effects could stem from the *antecedent* of the reciprocal (e.g. the subject *the men* in (1) or the nominal *foot-wide planks* in (2)), from the relational expression (e.g. *hit* or *stacked* respectively), or from the reciprocal expression itself. Of course, the combined interpretation of the three expressions can also lead to partitioning effects. In this paper, studying the availability of partitioned readings with different antecedents, we show new evidence that partitioned interpretations of simple reciprocal sentences are independent of the semantics of the reciprocal expression, and are exclusively determined by the interpretation of the antecedent. We propose that *definite* (as well as *indefinite*) antecedents can be interpreted as dependent on an implicit, contextually driven, quantifier, which triggers their partition. This factor was previously argued in Winter (2000) to be a crucial element for the analysis of other phenomena involving plurality and quantification.

2. Previous proposals

In the literature about reciprocals there is disagreement concerning the origin of partitioned interpretations. For sentence (1), many works (see Schwarzschild 1996; Dalrymple et al. 1998; Beck 2001) agree that the partitioned interpretation arises from a partitioning of the subject NP. The subject *the men* is assumed to be partitioned into two disjoint sets by a mechanism independently motivated for plural NPs. For each of the two sets, the interpretation of the reciprocal expression is determined independently of the partitioning of the NP denotation. In all of these works, the operator that is assumed to create this NP partitioning is the *cover* mechanism suggested in Schwarzschild (1996). This mechanism distributes a set denoted by a plural NP into contextually salient subsets, such that the union of the subsets equals the original set.

There is less agreement about the origin of the partitioned interpretation in cases like sentence (2) above. Dalrymple et al. (1998) (henceforth DKKMP) propose a system for the semantics of reciprocal expressions based on the principle called the *Strongest Meaning Hypothesis* (SMH). DKKMP's system includes a variety of available meanings,

and in each reciprocal sentence the logically strongest meaning that is consistent with relevant contextual information is chosen as the interpretation of that sentence. The partitioned interpretation of (2) is then derived by assuming that stronger, 'unpartitioned' meanings in their system are precluded because it is impossible for 'scores' of planks to form a single stack. The SMH accordingly selects a semantic operator called *Inclusive Alternative Ordering* (IAO) as the interpretation for the reciprocal in this sentence. The definition of the IAO operator allows the partitioned interpretation of sentence (2), as solely derived by the meaning of the reciprocal expression *each other*. Thus, DKKMP make a distinction between the origins of partitioning in (2), which is the SMH, and partitioning in (1), which is external to the meaning of the reciprocal expression. Another view on this kind of examples is offered in Beck (2001), where all partitioned interpretations are attributed to a general semantic process with plurals, using the cover mechanism of Schwarzschild (1996). In Beck's system, the IAO operator is not generated as one of the possible meanings of reciprocal expressions.

3. New evidence about partitioned reciprocals

The following minimal variation pair poses a challenge for both DKKMP's and Beck's analyses:

- (3) The planks are stacked atop each other.
- (4) Planks 1, 2, 3, and 4 are stacked atop each other.

Suppose there are two stacks of two planks each. Sentence (3) is true, whereas sentence (4) is false or very odd in this situation. This minimal pair shows that the type of the subject NP affects the availability of the partitioned interpretation: changing it from a definite plural NP to a proper name conjunction, without changing its denotation, eliminates the partitioned interpretation. The SMH mechanism cannot account for the contrast between sentences (2) and (3), in which partitioning is available, and sentence (4), in which it is not: in both cases forming one stack with four planks is easily possible, hence DKKMP's analysis expects both sentences not to show any partitioning effect. Schwarzschild's cover mechanism is also problematic, as it assumes no inherent difference between partitioning effects of different types of plural NPs: in both (3) and (4) the cover mechanism should allow the same partitioning effects, at least in the absence of explicit assumptions about its interaction with contextual factors.²

2. See Winter (2000) and Beck and Sauerland (2001) for discussion of such possible factors. In the case of sentence (4) above, we were not able to find contexts that clearly allow the partitioned interpretation, and similarly for the other reciprocal sentences below with proper noun conjunctions.

The effect of the type of the antecedent NP on the interpretation of reciprocal sentences is clearly exemplified when world knowledge only allows a partitioned interpretation. Consider the following sentences, in a situation where there are four singers:

- (5) The singers are looking into each other's eyes in this photo.
- (6) #John, Paul, George and Ringo are looking into each other's eyes in this photo.

Sentence (5) is felicitous, whereas sentence (6) is rather weird. In (6), despite world knowledge, the truth conditions derived from the reciprocal expression are not easily weakened to allow a partitioned interpretation.³

4. A new account of partitioned interpretations of reciprocals

In the literature on plurality, it has often been proposed that the way plurals are interpreted is governed by contextual factors that determine partitioning of their denotations. These factors are often independent of reciprocal expressions, and may therefore point to a general mechanism of partitioning. Beck's account that was mentioned above proposes to exploit the *cover* mechanism of Schwarzschild (1996) to capture such effects. Consider one of Schwarzschild's simple illustrations of partitioning, involving sentences like *The vegetables are too heavy to carry*. Clearly, the heavy objects referred to in this sentence do not have to be individual vegetables. They can also be baskets or boxes of vegetables. We adopt Schwarzschild's and Beck's view that such partitioning effects are strongly related to partitioned interpretations of reciprocals. However, we consider contrasts like the ones between sentences (3) and (4) as further evidence for the relations suggested in Winter (2000) between "partitioning" effects with plurals and the dependent/anaphoric interpretation of definite and indefinite descriptions. Winter shows that while definite plural NPs easily allow distribution to contextually salient subsets of their set reference, conjunctions of proper names often resist such distribution. The following example, adapted from Winter (2000), exemplifies this contrast.

- (7) The committee will commission operas to be written by teams of two composers.
 - a. The composers will earn \$100,000.
 - b. Lloyd Webber, Penderecki, and Stockhausen will earn \$100,000.

3. Note that a partitioned interpretation *is* available if the partition is syntactically expressed in the conjunction, as in the following variation of sentence (6) in (i) below. In this case, a partition to two pairs of singers is perfectly possible, as expected by compositionality and intersective ("Boolean") analysis of the italicized *and*.

(i) [John and Paul] *and* [George and Ringo] are looking into each other's eyes.

Consider a case where an opera was commissioned by the committee to be written by Lloyd Webber and Penderecki, while another opera was commissioned to be written by Lloyd Webber and Stockhausen. Each pair of composers received a total pay of \$100,000 for their opera. In this situation sentence (7a) is judged to be true whereas sentence (7b) is judged to be false, or very odd. According to Winter (2000), partitioning is available for the definite NP in (7a) because of the anaphoric power of the definite, which can combine with implicit quantification to create distribution into subsets. Informally, sentence (7a) is analyzed as follows:

- (8) For every team x , the composers in x will earn \$100,000.

This kind of *dependent* interpretation of definites is highlighted in the following example from Winter (2000) (cf. Partee 1989):

- (9) At a shooting range, each soldier was assigned a different target and had to shoot at it. At the end of the shooting we discovered that *every soldier hit the target*.

In the italicized sentence, the noun phrase *the target* is interpreted as a bound anaphor, dependent on the subject quantifier. The same mechanism does not operate on proper name conjunctions as in (7b) since they are not anaphoric.

This analysis captures the contrasts in (3)–(4) and (5)–(6), and more generally, it gives an alternative account of partitioning effects with reciprocal expressions and definite antecedents. For instance, for (3) we informally assume the following analysis.

- (10) In each group of planks g , the planks in g are stacked on top of each other.

A more complicated analysis may also work with DKKMP's example (2), provided we allow the quantifier over groups of planks, or situations, be in the scope of the subject, as in the following paraphrase:

- (11) He and scores of other inmates slept on foot-wide planks G s.t. in each subgroup g of G , the planks in g were stacked atop each-other.

This analysis agrees with the general assumption of previous works that partitioning is sensitive to context: the actual grouping of planks in the analysis (10) is determined by the context. However, crucially, this analysis of (3) relies on the anaphoric potential of the definite subject (cf. (7a) and (9)). In (4), with proper name conjunction, by contrast, partitioning is expected to be impossible due to the fact that the subject cannot be referentially dependent on any implicit quantifier.

A critical assumption in this analysis is that the interpretation of the reciprocal expression itself does not allow partitioning. For if partitioning could originate from the reciprocal itself, sentences like (4), with conjunctive antecedents, would also have been expected to show a partitioned interpretation. This suggests that there is a “lower bound” on the meaning of the reciprocal expression itself: it cannot be weak enough to

allow partitioned readings. We contend that all the cases of reciprocal sentences with partitioned interpretations are the result of an independent partitioning mechanism, while the reciprocal expression itself always has an unpartitioned reading. In more formal terms, we adopt the following assumption on the interpretation of reciprocals:

- (12) *Connectivity*: Let A be a set and let R be a binary relation. If the one-place predicate R *each other* holds of A , then the graph induced by R on A is *connected* (= not partitioned).

5. Further evidence

Once we examine further previously suggested reciprocal meanings that allow for partitioned interpretations, we see that these partitions as well are affected by the identity of the antecedent, similarly to the examples discussed above. Sentence (13) below is brought in DKKMP as an example for a reciprocal reading they call *One-way Weak Reciprocity* (OWR), which requires that each member of the antecedent set participate in the denoted relation with another member of the antecedent set.

- (13) “The captain!” said the pirates, staring at each other in surprise.

In sentence (13), OWR requires that each pirate stare at another pirate. This seems correct, for sentence (13) is true in both figures (2a) and (2b) below. However, consider what happens when we replace the subject in (13) by a proper name conjunction, as in sentence (14) below.

- (14) Morty, Charley, Oswald and Bob are staring at each other.

Here the truth conditions generated by OWR become less adequate, for sentence (14) is odder than (13) in Figure 2a, with a partitioned interpretation, while both sentences are perfectly OK in Figure 2b, with a non-partitioned interpretation. From this contrast between (13) and (14) we conclude that the actual interpretation of the reciprocal expression with the predicate *stare at* requires connectivity (i.e. lack of partitions) on top of the truth conditions required by OWR. The acceptability of (13) in Figure 2a is again attributed, as in (3) and (5), to the anaphoric potential of definites, independently of the meaning of reciprocal expressions.




Figure 2. pirates staring at each other

6. Conclusions

We have shown some evidence for a systematic contrast between the interpretation of reciprocal sentences with different plural antecedents. While reciprocals with plural definite antecedents often allow partitioning effects, reciprocal sentences with proper name conjunctions resist such partitioning. As in Winter (2000), we attribute this difference to the anaphoric potential of definite noun phrases – possibly on implicit quantifiers contributed by the context. Accounting for the main contrast we analyzed requires that the meaning of reciprocal expressions does not involve any partitioning effects. Thus, we believe that this paper has shown evidence that is pertinent for the theory of reciprocal expressions, as well as for the more general theory of plurality of quantification.

References

- Beck, S. 2001. Reciprocals are definites. *Natural Language Semantics* 9: 69–138.
- Beck, S. & Sauerland, U. 2001. Cumulation is needed: A reply to Winter (2000). *Natural Language Semantics* 8: 349–371.
- Dalrymple, M., Kanazawa, M., Kim, Y., Mchombo, S. & Peters, S. 1998. Reciprocal expressions and the concept of reciprocity. *Linguistics and Philosophy* 21: 159–210.
- Fiengo, R. & Lasnik, H. 1973. The logical structure of reciprocal sentences in English. *Foundations of Language* 9: 447–468.
- Partee, B. 1989. Binding implicit variables in quantified contexts. In *Papers from the 25th Regional Meeting of the Chicago Linguistic Society, CLS 25*: 342–65.
- Schwarzschild, R. 1996. *Pluralities*. Dordrecht: Kluwer.
- Winter, Y. 2000. Distributivity and dependency. *Natural Language Semantics* 8: 27–69.

The representation and processing of fixed and compositional expressions

Joost Schilperoord & Rein Cozijn*
Tilburg University

This paper examines the way language users process so-called regular fixed expressions. Such expressions are compositional in the sense that their overall meaning can be derived from their parts, while at the same time factors such as usage frequency, resistance to modification, ‘nativelikeness’ suggest that they are stored and processed as complex lexical units, just like idioms. To test this hypothesis, we experimentally investigated the availability of regular fixed expressions for anaphoric reference as compared to fully compositional expressions. If an expression is processed as a unit, its parts will be less available, and therefore, processing anaphoric expressions should be more difficult and more time consuming compared to fully compositional expressions. The results supported this hypothesis.

1. Introduction: Storage versus computation

Cognitive theories of language and language processing usually draw a distinction between those parts of linguistic knowledge that speakers already know at the time of use, and those parts that are independent from prior language experience, i.e., that speakers find out *on the fly*. Language processing consists of a delicate balance between what is stored in the mental lexicon and what is constructed by the mental grammar. The usefulness of this distinction is evenly entrenched and acknowledged. However, if it comes to matters of *distribution* (what *is* stored and what *is* computed?), views begin to diverge considerably. In his *Foundations of Language*, Jackendoff puts the issue at stake as follows:

“What aspects of an utterance *must* be stored in long-term memory, and what aspects *can* be constructed online in working memory?” (Jackendoff 2002: 152).

*We would like to thank Marleen Smits for conducting the experiment reported in this study. Thanks are also due to two anonymous reviewers of this paper.

The answer that ‘mainstream’ theories of language provide is that the memorized/stored units are the *words* of the language (roughly equivalent to X^0 -categories): those elements that are devoid of the kind of combinatorial structures that can be described by the grammar. In other words, the line between memorization and computation is drawn between what can and what cannot be predicted by the rules. We shall provisionally refer to this view as the ‘words-only’ hypothesis. This hypothesis represents the *strongest* possible one regarding the distinction between lexicon and grammar and is generally considered the best account of what (generative) linguists take to be *the* outstanding property of language: Discrete infinity; the fact that speakers are able to produce an infinite set of sentences from a finite grammar – “(...) virtually every sentence that a person utters or understands is a brand-new combination of words, appearing for the first time in the history of the universe” (Pinker 1994: 22). However, over the past decade cognitive linguists and psycholinguists alike have severely challenged the words-only view (see, for example, Jackendoff 1997, 2002; Goldberg 1995; Pinker & Jackendoff 2005; Langacker 2000)¹ and have come to favor a more relaxed view on the distinction between lexicon and grammar to the effect that speakers memorize a lot more than the words-only hypothesis predicts. People’s everyday speech shows that many, if not most, utterances consist of, or contain somehow conventionalized, preformatted expressions like speech formulas (‘How are you?’ ‘Thanks, I’m fine’), sayings (‘liever lui dan moe’, *rather lazy than tired*), idioms (‘hij ging de pijp uit’, *he kicked the bucket*) or proverbs (‘home sweet home’). Such expressions are characterized by a restricted selection of words, to be uttered in a more or less fixed order. They furthermore display quite specialized aspects of meaning and often resist modification. This has lent support to the idea that such fixed, multi-word combinations are stored as lexical items as well, and are retrieved from memory at the time of use rather than being computed from ‘scratch’. Although the literature shows a bewildering number of terms to refer to such fixed combinations,² one generally accepted term is that of *fixed expressions* (or FEs). Rather than discrete infinity, cognitive linguists have come to acknowledge an alternative, possibly even more compelling fact about the cognitive representation of language: Speakers appear to know ‘by heart’ literally thousands of FEs.

1. Earlier testimonies can be found, for example, in Bolinger (1976) and Chafe (1968). Current interest in fixed expressions owes a lot to sociolinguists, who have acknowledged the importance of formulaic language already for a long time.

2. Wray and Perkins (2000) mention over 40 (!) terms. They can be form-based (‘polywords’), based on semantic irregularity (‘idioms’), or based on resistance against regular syntactic operations (‘fossils’) or on formulaicity (‘sayings’, ‘vocabulary’). Yet other terms are based on usage frequency or usage function.

Recent attempts to incorporate FEs of various types into models of linguistic knowledge and linguistic processing converge on what might be labeled *hybrid* models of language use – models that combine rule based computation on the one hand, and massive, storage based memorization on the other. The combinatorial principles of the grammar allow speakers to assemble individual lexical items into ‘novel’ utterances (the ones appearing ‘for the first time in the history of the universe’) or to comprehend them, whereas preformatted strings of words can be used and brought forth whenever cognitive and communicative circumstances license their usage. Proposals (of various kinds and with varying levels of detail) for hybrid models can be found, for example, in Langacker (1988), Kuiper (1996), Schilperoord (1996), Wray and Perkins (2000), Jackendoff (2002), Sprenger (2003), Dabrowska (2004), and Jackendoff and Pinker (2005). Hence, such models are based on a general distinction between combinatorial and configurative structures and processes, thus predicting that FEs constitute a *general property* of linguistic knowledge, rather than some exceptional area of linguistic exotica.

The empirical research on FEs that fuels the development of such models has, for a long time, been confined to the study of *idioms* (but see Kuiper et al. 2007). Especially idioms present a challenge to words-only models. Idioms are ‘complex units’ (Everaert et al. 1995: 3) – *complex* in that they contain more than one word, and *units* because of their more or less stipulated meanings: One cannot derive the meaning of an idiom from the meanings of its constituting parts. For example, the Dutch expression *de plaat poetsen* (‘to sweep the plate’) means ‘to leave’, but if a speaker would look up the meanings of the distinct elements *plaat* and *poetsen* and combine these according to a syntactic rule V + DO, s/he would not be able to arrive at any other conclusion than that – literally – a plate is swept. Other idioms show, at least to a certain degree, compositionality, but the fact remains that their overall meaning is ‘specialized’ with respect to the meanings of their constituting parts, and cannot be derived via the grammar of the language concerned.

Another reason why idioms challenge word-only models of processing is that, often, they display unusual patterns of syntactic modifiability. If, for example, the expression *de plaat poetsen* is used in a passive construction like *de plaat werd geпоetst* (the plate was being swept, see Van Gestel 1995), it no longer carries the idiomatic meaning *to leave*. Hence, idioms are paradigmatic cases of lexically listed items ‘larger than the word’: Various atomic units (X^0 elements) are involved, and yet the correspondence between these units and the overall meaning concerns some fixed combination, rather than rule-based composition.

Apart from such typical linguistic topics in the study of FEs, psycholinguists are especially interested in the ways language users *process* idioms (see, for example, Sprenger 2003; Van de Voort & Vonk 1995; Tabossi & Zardon 1995). Research issues

concern, among others, the way language users extract idiomatic meanings from linguistic expressions – either by mere retrieval or by online construction, and the role that the literal meaning of the lexical elements of idioms play during online comprehension (Sprenger 2003, and references therein).

Despite their significance for the study of FEs, idioms are certainly not the only type of FEs. In fact, in this study we will focus on a broad class of linguistic expressions that, at first sight, appear to be entirely compositional. To set the stage, consider the examples in (1).

- (1) a. weapons of mass destruction
- b. *het puntje van m'n tong*
 'the tip of my tongue'
- c. (*iemand, iets*) *het voordeel van de twijfel* *geven*
 'to give (someone/something) the benefit of the doubt'
- d. *naar aanleiding van uw brief/email/telefoontje....*
 'in reference to your letter/email/telephone call....'

Expressions like these can be considered compositional in that anyone who knows the meaning of, for example, the words *geven*, *voordeel* and *twijfel* (1c) and the relevant phrase structure rules of Dutch, is able to find out what *het voordeel van de twijfel geven* means. Hence, unlike an idiom, it would seem that the conceptual interpretation of the expressions in (1) follows, at least partly, from the 'normal' rules of combining the meanings of their parts, which seem to mean the same as they do when used in isolation. Yet, intuition tells us that expressions like *het voordeel van de twijfel* or *het puntje van m'n tong* are preformatted units as well, just like idioms, and that they belong to the repertoire of stored lexical items that speakers may bring forth 'en bloc'.³ What motivates such an assumption? In fact, this question motivates the current research. Our aim is to bring online processing evidence to bear on the issue, evidence that shows that expressions like (1) are in fact processed differently from entirely 'novel' (i.e., fully compositional) phrases and sentences, and that these differences testify to the idea that their actual usage and understanding is based on *recognition* rather than on compositional analysis by the language grammar.

For now, however, let us briefly mention some informal observations that suggest storage. For one thing, *recognition* is an important factor testifying to the configurative (stored) nature of expressions like (1a–d). Native speakers judge them to be 'familiar',

3. Example (1)d shows that preformatted expressions need not be entirely fixed. The expression *naar aanleiding van uw....* contains one variable slot. Possibilities for inserting free elements are limited, however. In this particular expression, the slot has to contain some communicative artefact, like a phone call, an advertisement, a letter, and so on.

that is, they recognize them as “(...) being the way native speakers of their language say that sort of thing” (Kuiper 1996: 3). Hence, (1a–d) (and comparable expressions) define native competence, and their proper use is part and parcel of our ability “(...) to sound nativelike” (ibid. 3). In more controlled experimental settings, native speakers have been shown to be perfectly able to restore the original expression when presented to them with words omitted or scrambled (see Schilperoord & Peeters 1998).

Another factor suggesting memorization rather than computation of (1a–d) is their *resistance to modification*. As an example, consider once again the expression *het voordeel van de twijfel*. It is used in a situation in which some decision has to be made that may turn out to be positive or negative with respect to the person or situation being evaluated. The expression itself then refers to a *positive* outcome. However, if the speaker, using the expression, also wants to communicate the idea that the decision that was made was not controversial, s/he may modify the second noun, for example by expressing *het voordeel van de lichte twijfel* (the benefit of the slight doubt). The adjective *slight* thus refers to the nature of the decision to be made. Although this variant is not ill-formed in any way, and perfectly understandable, speakers turn out to prefer the (1c) variant, even over this slightly modified alternative. We performed a Google search that produced over 124.000 hits for *het voordeel van de twijfel*, as opposed to only two hits for the variant expressions *het voordeel van de lichte twijfel*.⁴

Although such informal observations strongly suggest lexical listing, the way the lexical primitives of, say, *het voordeel van de twijfel* relate to its overall meaning differs from that of idioms. If a speaker doesn't know the meaning of *de plaat poetsen*, there is no rule that will tell him that, whereas rule based analysis should suffice to find out what *het voordeel van de twijfel* means. On the other hand, especially their resistance to alteration suggests some kind of semantic specification as well. Because expressions like (1) thus seem to combine conceptual transparency and pre-formatting, we shall refer to them, perhaps somewhat awkwardly, as *regular fixed expressions*, or RFEs.⁵

As already indicated, our goal here is to turn the issue of RFEs as possibly stored complex units into an empirical affair, and thereby to contribute to the ongoing debate on the storage versus computation issue. In the next section, we report the set up and

4. Search date: 11/29/2007. The expression also resists pluralizing one of the nouns, or replacing the definite article by an indefinite article. These factors, too, testify to lexical listing.

5. Mel'čuk (1995) calls these *non-free phrases*, or *phrasemes*. One reviewer suggested that 'our' RFEs may be *compositional restricted collocations* (cf. Kuiper et al. 2007). Given the fact that *het voordeel van de twijfel* resists replacing the definite articles by indefinites, this may actually be the case. However, we're not sure whether this is the case for all experimental expression that we used (see appendix). The issue does however testify to the idea that the notion of 'fixed expression' suggests a far more homogeneous class than is actually the case. We leave this issue for further research.

results of an experiment in which we compared the processing of such expressions with their non-fixed counterparts. The experiment is based on the following processing characteristic: If an expression is stored, and retrieved from the lexicon ‘en bloc’, little attention will be devoted to its parts as compared to fully compositional expressions (that have to be processed in word-by-word fashion). If this can be experimentally shown to be the case, such a finding would be evidence in favor of the storage assumption. How we have attempted to accomplish this, is described in Section 2. In the final section, we will return to the storage issue and elaborate on a theoretical underpinning of the way speakers mentally represent their knowledge of RFEs like *het voordeel van de twijfel*. This section draws on the results of the experiment and on current theorizing on the lexicon and grammar. It provides a more thorough and detailed account of the informal observations described in the introductory section with regard to the stored nature of RFEs.

2. The experiment

The basic idea behind the experiment is to compare processing aspects of RFEs and fully compositional expressions. To explain how we have attempted to accomplish this, we shall discuss in some detail one of the experimental stimuli. During the experiment, participants read either (2a) or (2b).

- (2) a. *Met het bestuur heb ik het functioneren van u binnen onze firma besproken. Tijdens het overleg is ieders mening gehoord. Wij willen u het voordeel van de twijfel geven in deze bijzondere situatie. Zoals u weet werd die opgeroepen door uw matige prestaties.*
 ‘I have discussed your performance in our company with the board. During that meeting, everybody’s opinion has been taken into consideration. We want to offer you the benefit of the doubt in these specific circumstances. As you know that was raised due to your poor performance...’
- b. (...) *Wij willen u het voordeel van de lichte twijfel geven in deze bijzondere situatie. Zoals u weet werd die opgeroepen door uw matige prestaties.*
 ((...)
 ‘We want to offer you the benefit of the slight doubt in these specific circumstances. As you know that was raised due to your poor performance...’

Example (2) presents two short discourses that contain either the RFE *het voordeel van de twijfel* (a) or a slightly modified version *het voordeel van de lichte twijfel* (b). The modification consists of the insertion of the adjective *lichte* (slight) before the second

noun of the construction, rendering the construction non-fixed. If in fact *het voordeel van de twijfel* is stored as a unit, processing it will result in the retrieval of its meaning which will be inserted 'en bloc' into the discourse representation. The processing of the modified version, on the other hand, should give rise to the kind of computation processes that characterize the processing of fully compositional expressions, resulting in a discourse representation in which each of the distinct parts of the modified RFE are integrated separately. Hence, the two parts of the expression, i.e., *voordeel* and *twijfel*, will be stored as one lexical unit and not represented as such. In case of (2b), on the other hand, the two parts will be represented individually.

The differential representation of the parts of the expression, we hypothesized, should have consequences for their availability to subsequent processing, for instance as antecedent of an anaphoric expression. The discourses in (2) contain such an anaphoric expression: The pronoun *die* in the sentence *Zoals u weet werd die opgeroepen door uw matige prestaties* refers back to the entity *twijfel* in the sentence with the (modified) RFE. Resolving the anaphora involves finding the antecedent in the discourse representation built up so far. Since the entity *twijfel* is not represented as such in the discourse representation of (2a), but *is* represented in the discourse representation of (2b), resolving *die* to *twijfel* should be easier in (2b) than in (2a).

Summarizing our line of reasoning, antecedents in RFEs are assumed to be less available than antecedents in fully compositional expressions. Processing anaphoric expressions should, therefore, be more difficult and more time consuming if the antecedent is part of a regular fixed expression than if it is part of a fully compositional expression.

This hypothesis was tested in a reading experiment where eye movements were monitored while participants read short texts containing an anaphoric expression that pointed to an entity in either an RFE or in a modified RFE. The reading times of the region with the anaphoric expression as well as the number of regressions from that region to the antecedent were registered. When readers run into difficulties while processing text, they typically engage in three strategies: They either stay longer in the region where the difficulty arose, regress to an earlier part of the text where they expect to find information that is helpful in resolving the issue, or proceed to the end of the sentence and re-read the sentence entirely (Frazier & Rayner 1982). Since resolving anaphoric expressions involves finding a referent earlier in the text, it was expected that readers would apply the first and the second strategy, but not the third. This means that longer reading times in the sentences containing the anaphoric expression or more regressions from that sentence to the region containing the antecedent should occur if the antecedent is part of an RFE than if it is part of a modified RFE.

Two types of reading time measures were calculated from the eye movement data: *forward reading time* and *total-pass reading time*. *Forward reading time* represents the time spent in a region from the start of the first fixation in that region, provided that the reader enters that region for the first time, until the end of the last fixation in that

region, including possible regressions within the region, before the region is left in a forward direction (Vonk & Cozijn 2003; Murray 2000). The measure captures reading patterns associated with the first strategy mentioned above. *Total-pass reading time* is an extension of the *forward reading time* and includes the time readers spend on earlier regions of the text. It therefore reflects the second strategy mentioned above. Of course, a simple count of the number of regressions will be indicative of the second strategy too.

Furthermore, the reading times of the RFEs and the modified RFEs were registered as well. It is unclear, however, whether the assumed, differential processing of RFEs and modified RFEs will result in differences in reading time. On the one hand, it might be assumed that RFEs will be processed faster, since less attention will be paid to their constituting parts. On the other hand, the processing of RFEs will result in the retrieval of their unified meaning which might consume time. Analysis of the reading times might shed light on this issue.

The reading experiment entailed the reading of several texts that either contained an RFE or a modified RFE and verifying statements about the text. The experiment will be explained below.

2.1 Method

Participants

A total of 43 students of Tilburg University participated in the experiment. The data of ten participants could not be used in the analyses: Three participants had made too many errors in the verification task, two participants had not followed the instructions, and five participants were excluded because of measurement errors. In all, the data of 33 participants, 19 female and 14 male, ranging in age from 19 to 26 years (21.6 on average) were submitted to analysis.

Materials

Twenty-four fixed expressions like *voordeel van de twijfel* were obtained from newspapers, (in)formal letters, conversations on radio or television, and De Coster (2002). RFEs are hard to identify. Erman and Warren (2000) specified two criteria for their identification: *restricted exchangeability* and *favourability*. The first criterion, stating that “at least one member of the prefab [prefabricated expression] cannot be replaced by a synonymous item without causing a change in its meaning, function and/or idiomaticity” (p. 32), is hard to apply because it requires a subjective interpretation on the part of the analyst. The second criterion, however, is more tangible. It states that fixed expressions ‘are favored by native speakers in preference to an alternative combination which would have been equivalent had there been no conventionalization’ (p. 31). So, the status of an RFE can be established on the basis of judgments by native speakers. This method was used to evaluate the fixed expressions.

The RFEs were embedded in short texts that supplied a suitable context for the expressions to occur. The texts consisted of six sentences: two introductory sentences, a sentence with the RFE, a sentence with an anaphoric expression referring back to an entity in the RFE, and two concluding sentences. For each text, a version was created that contained a modified form of the RFE. The modification consisted of the insertion of an adjective into the expression, rendering the expression non-fixed. A Google survey on the internet showed that the ratio of occurrences of the RFEs to the modified RFEs was (minimally) 3000:1.

The 24 texts were submitted to a judgment study in which the naturalness of the texts as well as the clichédness of the expressions was tested. The 24 texts were split up into two lists of 12 texts. Each list contained six unmodified and six modified text versions, which were distributed in a semi-random fashion. Of each of the two lists a counterpart was created that contained the same texts but in the mirrored condition. This resulted in four lists in total. Forty participants, undergraduate students of Tilburg University, were assigned randomly to one of the four lists, so each list was handed out to 10 participants. The participants were instructed to read the texts and to evaluate the connectedness of each sentence with the prior text. If the participants felt that a sentence did not connect well, they should underline the sentence and specify a reason below the text. They were free to write down other remarks about the texts as well. After the reading task, the participants were presented with a list of the 24 fixed expressions used in the texts. They were asked to indicate on a scale ranging from 1 (not very) to 7 (very) how clichéd they believed the expressions were. This second judgment task concluded the materials study. The complete test took about 20 minutes.

The results of the naturalness tests showed that two texts received many comments. They were excluded from the materials set. The (few) comments on the other texts were used to improve the texts. No striking differences were found between the modified and unmodified versions of the texts. The average clichédness score of the 24 fixed expressions was 4.33 (in a range of 2.83 to 5.55). An analysis of variance was performed on the clichédness scores to see whether having seen a (modified) RFE embedded in a text in the naturalness test prior to the clichédness test had influenced the judgments. This was not the case (all $F_s < 1$). The four RFEs with the lowest scores (below 3.75) were removed from the materials set. The complete materials test resulted in 18 experimental texts with an average clichédness score of 4.61 that contained an RFE in two conditions: with or without a modifying adjective.

As mentioned above, the fourth sentence of each text contained an anaphoric expression that had to be resolved to an entity that was part of the RFE in the third sentence. The type of anaphoric expression was varied: It could be singular or plural, male/female or neutral, and pronominal or demonstrative. The sentences with the RFE always contained exactly two entities that could serve as antecedent of the anaphoric expression, rendering the anaphoric expression ambiguous. The distance between the

anaphoric expression and the intended antecedent was kept constant in all texts and measured nine intervening words.

For each experimental text a verification statement was constructed that required the anaphoric relation in the text to be resolved. For instance, the verification statement of the text in (2) read: "*De twijfel werd opgeroepen door de matige prestaties*" (*The doubt was raised due to the poor performance*). Correct verification implied that the readers had connected the anaphoric expression *die* (*that*) in the fourth sentence to the entity *twijfel* (*doubt*) in the third. The task ascertained that the texts were read thoroughly. For a complete list of all experimental texts, see the appendix.

To conceal the goal of the experiment, the 18 experimental texts were complemented with 18 filler texts that resembled the experimental texts in topic and style. Furthermore, the filler texts varied in length from five to eight sentences and had verification statements that were untrue, thus balancing the 18 true statements of the experimental texts. The verification statements of the filler texts related to information in other parts of the text than the third and the fourth sentences, as in the experimental texts. Finally, three texts were created that served as practice items in the experiment.

Design

The experimental and filler texts were distributed semi-randomly over a list of 36 texts, with no more than two texts of the same type in a row. Within the list, half of the experimental texts contained the unmodified version and half the modified version of the RFE. The distribution of the two versions was semi-random as well. The 18 experimental and 18 filler texts were assigned to the 36 slots in the list such that texts that showed a resemblance in topic were positioned as far apart as possible.

Since each version of an experimental text could be presented to a participant only once, a second list was created that was identical to the first but that mirrored the first in experimental condition. If a text was without modification on the first list, it was with a modification on the second, and vice versa. Participants were assigned randomly to one of the two lists.

Procedure

The experiment was conducted in the research lab of the Department of Communication and Information Sciences at Tilburg University, using an SR Research EyeLink 2 eye-tracker system, operating at 250 Hz. Measurements were performed on the right eye only.

Participants were told that the study's aim was to increase our understanding of how people read texts. They were instructed to read several texts and to verify statements about them afterwards. After the instruction, the installation of the eye-tracker equipment, and its calibration, the first trial appeared on the computer screen. Each trial started with a recalibration point on the screen at the position where the first

word of the text would appear. The recalibration procedure ascertained proper measurement as well as fixation of the eyes at the starting position of the text. By pressing the enter-key on the keyboard the recalibration point disappeared and the text appeared. The participants had to read the text as they normally would and press the key as soon as they had reached the end of the text. Immediately after pressing the key, the text disappeared and a verification statement appeared at the position of the last sentence of the text. The statement had to be verified, as fast as possible, in relation to the text that had just been read. Judgments were to be given by pressing either the right most key (true) or the left most key (false) of the lowest row on the keyboard. After the judgment had been given, the statement disappeared and the recalibration point of the next text appeared indicating the start of the next trial. The complete experiment lasted approximately 30 minutes.

Preparation of the data

The eye-movement data were analysed using the program *Fixation* (Cozijn 2006). The fixations were assigned automatically to the words in the text and then checked manually. In order to be able to analyse the processing of the RFEs and the anaphoric expressions, their embedding sentences were divided into regions as in (3). The regions are indicated by numbered slashes.

- (3) *We willen u*^{1/} *het voordeel van de*^{2/} *lichte*^{3/} *twijfel*^{4/} *geven*^{5/} *in deze bijzondere situatie.*^{6/} *Zoals u weet werd*^{7/} *die*^{8/} *opgeroepen door uw matige prestaties.*^{9/}
 We want you^{1/} the benefit of the^{2/} slight^{3/} doubt^{4/} give^{5/} in these specific circumstances.^{6/} As you know was^{7/} that^{8/} raised due to your poor performance.^{9/}

The regions were defined as follows: (1) the beginning of the RFE sentence, (2) the part of the RFE up to the antecedent entity or the modifying adjective (if present), (3) the modifying adjective, (4) the antecedent entity of the RFE, (5) the final part of the RFE, (6) the final part of the RFE sentence, (7) the part of the next sentence up to the anaphoric expression, (8) the anaphoric expression, and (9) the final part of the sentence containing the anaphoric expression. Not all texts contained all regions. In the versions with an unmodified RFE, region 3 was absent. In thirteen texts, region 1 was absent. This was the case if the sentence started with the first part of the fixed expression. One of these texts did not contain region 2, but started immediately with the antecedent entity or the modifying adjective (if present). Twelve texts did not have region 5, because the RFE ended with the antecedent entity. Finally, in five texts the sentence with the anaphoric expression was concluded by a main or a subordinate clause after region 9. This extra clause was numbered 10. The regions that were consistent over texts were 4, 6, 7, 8, and 9.

The analysis program calculated for each region the *forward reading time*, the *total-pass reading time*, and the number of regressions.

2.2 Results

The data of two texts were excluded from analysis. Text 3 was excluded because it appeared to have been re-read several times by several participants due to some kind of reading difficulty. Text 14 was excluded because too many errors were made on its verification statement (77%). Furthermore, the eye-movement data of four different texts read by four different participants were excluded due to measurement errors (.8%). Of the remaining data, all those cases were excluded where an error had been made on the verification statement (10.0%). Finally, if reading times exceeded two standard deviations from the participant and the item means in a condition, they were considered outliers and excluded as well. Given these restrictions, the data of 33 participants and 16 texts were entered into the analyses.

For the *total-pass reading times* and the *forward reading times* in each region as well as for the verification times, analyses of variance were performed by participants (*F1*) and by items (*F2*) with Modification as a within participants and within items factor. The between-participants factor List was included in the participants analyses, solely for purposes of error variance reduction (Pollatsek & Well 1995). It will not be discussed below. The following regions were analysed: region 2+4+5 (the RFE), region 4 (the antecedent), region 6 (the final part of the sentence containing the RFE), region 8 (the anaphoric expression), region 9 (the part of the main sentence following the anaphoric expression), and region 8+9.

The number of regressions from regions 8+9 and 8+9+10 (the anaphoric expression and the following text) to regions 4 (the antecedent) and 2+4+5 (RFE) were analysed with χ^2 -tests.

Reading times and verification times

The mean reading times and verification times are presented in Table 1.

Table 1. Mean *total-pass reading times* and *forward reading times* (both in ms) of the antecedent (2+4+5, 4, and 6) and the anaphoric regions (8, 9, 8+9), and the mean verification times (ms) as a function of Modification of the RFE (percentages of regressions are in parentheses). The final columns show the effects (*F1*, *F2*) and the percentages of valid observations:

Measure	Region	Modification		<i>F1</i>	<i>F2</i>	%Obs
		Unmodified	Modified			
<i>Total pass reading times</i>	2+4+5	1008(11.3%)	937(13.2%)	†		89.3
	4	436(23.3%)	392(17.7%)	†		74.2
	6	1112(11.8%)	1147(18.6%)			89.7
	8	317(13.0%)	273(9.6%)			55.7

(Continued)

Table 1. Continued

	9	1574(40.3%)	1396(31.8%)	*	†	89.5
	8+9	1818	1574	*	*	89.5
<i>Forward reading times</i>	2+4+5	676	641			82.4
	4	376	322	*		58.0
	6	1054	1017			74.8
	8	241	249			49.2
	9	1081	1175			56.7
	8+9	1005	1086			71.6
Verification times		3335	3346			90.0

† $p < .10$; * $p < .05$

According to the total pass reading times, modified RFEs were read faster than unmodified RFEs (regions 4 and 2+4+5). However, this was only a trend in the *F1* analysis and there was no effect in the *F2* analysis. Similarly, the *forward reading times* of region 4 showed that the region was read faster in the modified condition than in the unmodified condition. This was true for the *F1* analysis only. The reading times of the anaphoric region (9 and 8+9) revealed an effect of Modification: If the RFE had been modified, reading times were shorter than if the RFE had not been modified. There was no effect of Modification on the verification times.

Regressions

Table 1 shows the overall percentages of regressions for unmodified and modified RFEs. However, these percentages are not indicative of what region the regressions were directed at. A closer look at the regressions from the anaphoric region to the antecedent region showed that there were significantly fewer regressions if the RFE had been modified than if it had not been modified. The numbers of regressions are shown in Table 2.

Table 2. Number of regressions from the anaphoric regions (8+9 and 8+9+10) to the antecedent (4) and the RFE (2+4+5)

To region	From region	Modification	
		Unmodified	Modified
4	8+9	10	0
	8+9+10	39	4**
2+4+5	8+9	20	7*
	8+9+10	126	86**

* $p < .05$; ** $p < .01$

To summarize, the reading times of the part of the sentence following the anaphoric expression as well as the regressions from that part to the RFE showed that the resolution of the anaphoric expression caused more difficulties if the antecedent was part of an RFE than if it was part of a modified RFE, indicating that the entities contained in RFEs are less available than those in fully compositional expressions.

The results of the experiment, therefore, indicate support for a hybrid model of language use, a model that combines rule based computation with massive, storage based memorization.

3. Discussion

In this final section, we shall discuss the results of the experiment in the light of cognitive linguistic theories of the representation of linguistic knowledge. As we have already noted, to consider a string of lexical elements to be fixed can be motivated on various grounds: *usage based* factors such as frequency of use or sounding 'nativelike', and *cognitive* factors such as modifiability or conceptual specification. It will come as no surprise that our chief interest concerns cognitive factors. From a cognitive point of view, an expression can be considered 'free' if it exists of a combination of primitive units (say words) that, in principle, can be accounted for by the combinatorial rules of the language, i.e., (morpho)syntactic rules and conceptual rules. In such cases, the conceptual structure conveyed by the expression as a whole is entirely motivated by the meaning of its constituting elements and rules of combination. Hence, in a quasi formalized way, a free expression can be represented as (4a). What this means is that the meaning that units a/A and b/B contribute to the meaning of the whole utterance is the same as when used in isolation. This is indicated by the subscripts 1 and 2.⁶ Therefore, a/A and b/B are stored lexical primitives, whereas their combination is constructed online.

- (4) a. $(a)_1 (b)_2 / [A]_1 [B]_2$
 b. *de sleutels₁ van₃ de flat₂*
 the keys₁ of₃ the apartment₂ i.e. 'the keys to the apartment'
 c. $[SLEUTEL]_1 \text{RELATION}_3 [FLAT]_2$

6. Following a convention introduced by Jackendoff (1997), we use subscripts to express the form-meaning pairing of a certain unit. Each couple (a)/[A] with identical subscripts constitutes a lexical item. () represent *form* aspects, [] *meaning* aspects. Such relations are termed *correspondences* by Jackendoff. A lexical item, of whatever size, thus is a correspondence of phonological, syntactic and conceptual structure. Phonological structure is not presented here.

An example is the general phrasal template ($_{NP} (art) N_1 (_{PP} prep ((art) N_2))$), like in (4b), where the Dutch preposition *van* ('of') instantiates the preposition slot. This template can be used to express a conceptual relation between two nouns, or rather the objects conveyed by the two nouns. The relation itself may vary – it may express allowance (like in (4b)), ownership (e.g., *the house of the neighbors*), locative relations (e.g., *the centre of the field*), or part-whole relations (e.g., *the goal keeper of the team*) – but the relevant issue is that, whatever the relation, in a free expression like (4b), the nouns *sleutel* and *flat* mean exactly the same thing as when used in sentences like *ik ben m'n sleutels kwijt. Ik hoop dat ik m'n flat nog inkom vandaag* (*I've lost my keys. I hope I will be able to enter my apartment*) – hence the subscripts 1 and 2. The conceptual relation between the two nouns, as shown in (4c), can thus be fully attributed to the preposition (hence subscript 3). The compositional nature of (4) can be further substantiated in that it allows for 'normal' extraction of the prepositional object (see 5).⁷

- (5) *Dit is (de flat) waar ik (de sleutels (van (Ø)) heb.*
 This is (the apartment) which I have (the keys (to (Ø))
 'This is the apartment to which I have the keys'

Once again, in (5) the meaning of the utterance is composed of the meanings of its constituting parts. These considerations motivate a definition of a fixed expression as an expression, either conveying a conceptual structure that cannot be analyzed in terms of its constituent parts, or that exhibits syntactic inflexibility. Take the following Dutch expression (6).

- (6) *het neusje van de zalm*
 the little nose of the salmon → 'THE BEST'

The relevant observation is that in (6) the nouns do not contribute to the overall meaning the way they do if used in isolation. In addition, the expression does not express a relation between the two nouns. Note, however, that the expression instantiates the exact same syntactic phrase ($_{NP} (art) N_1 (_{PP} prep ((art) N_2))$) as does (4b). Although many have claimed that the elements of expressions like (6) can be distinctively motivated (see, for example, Van Gestel 1995; Gibbs 1995; Nicolas 1995; Geeraerts 1995), the fact remains that in (6), the nouns do not mean what they normally mean. Actually, (6) doesn't evoke the idea of either a salmon or a nose which is the very reason why it is an *idiom*. In principle, it allows for extraction of the prepositional object, but if that occurs, the expression loses its idiomatic ('figurative') meaning (see 7).

- (7) *Dit is de zalm waar ik het neusje van geproefd heb.*
 This is the salmon of which I the little nose have tasted
 'Of that salmon I have tasted the little nose'

7. See also Deane (1991).

The sentence in (7) does actually refer to a salmon and a nose, but it has lost its idiomatic meaning. Hence, an expression is fixed if its conceptual structure does not follow from the meanings of its constituents (Geeraerts 1995: 57 calls this ‘semantic specialization’). It can be represented in the general format given in (8).

- (8) ((a) (b))₁ / [c]₁

The subscript 1 indicates that the correspondence between the distinct elements on the one hand and the conceptual structure on the other is located at the level of the entire expression, rather than on the level of its constituting parts. This is why expressions that instantiate this template are to be considered lexical items by themselves. They are stored and retrieved at the time of use, rather than composed and constructed ‘on the fly’.

Idioms – expressions whose overall meaning cannot be derived from its parts – are paradigmatic cases of FEs. If one doesn’t know the meaning of *het neusje van de zalm*, there is no rule that will tell one so. But now, let us have a look at the expression in (9), once again instantiating the (_{NP} (art) N₁ (_{PP} prep ((art) N₂))) format.

- (9) *het voordeel van de twijfel*
‘the benefit of the doubt’

The expression in (9) is of the type that we used as stimuli in our experiment and which we called regular fixed expressions, RFEs. The expression appears to be compositional rather than configurative. Anyone who knows (the meaning of) the words *voordeel*, *van*, and *twijfel*, and the relevant phrase structure rules of Dutch, should be able to find out what (9) means. In addition, it would seem that the constituting lexical elements contribute to the overall meaning of the phrase in a fully predictable way. Hence, its representation would seem to follow the pattern in (4a) and not the pattern in (8). However, as we argued earlier and as the results of the experiment suggest, an expression like *het voordeel van de twijfel* appears to be preformatted and to belong to the repertoire of stored lexical items in that its actual usage and understanding is based on *recognition* rather than on analysis by the language grammar.

This can be further substantiated by the fact that, like (6), expression (9) does not allow extraction of the prepositional object, as shown in (10). Extraction renders the sentence unacceptable, if not plainly ill formed.

- (10) ??Het is *de twijfel* waar Jan *het voordeel van* heeft.
??It is *the benefit* which Jan has *the doubt of*

The impossibility of extraction suggests that the overall meaning of the expression in (9) is also specialized, albeit in a different way as, for example, *het neusje van de zalm*. In our view, the way the lexical elements in (9) relate to the overall meaning occupies an intermediate position between (4a) and (8) in that the overall meaning appears to be

integrated. Admittedly, this cannot be stated as rigidly formal as is the case with idioms like *het neusje van de zalm*, but the idea is that *het voordeel van de twijfel* expresses one unified conceptual structure⁸ – a *conceptual gestalt*, so to speak – instead of two related, but otherwise separate, concepts, while at the same time the elements making up this unified concept still appear to be recoverable. Another way of saying this, is that *het voordeel van de twijfel* doesn't refer to either *voordeel* (benefit) or *twijfel* (doubt). Hence, the representation of the expression may be given as (11).

$$(11) \quad ((a) (b))_1 / [A-B]_1$$

The structure in (11) is not entirely compositional – hence subscript 1 – which is the reason why (11) differs from (4a), where the relation (subscript 2) between two objects (subscripts 1 and 3) can be conceptually separated from the objects themselves. Note, however, that the two concepts making up the unified conceptual structure in (11) are both present, rendering the relation between the lexical elements and the overall conceptual structure still transparent, which accounts for the fact that the meaning of this expression is still analyzable. This is why (11) differs from (8).

From this analysis, it follows that the phrasal structure (_{NP} (art) N₁ (_{PP} prep ((art) N₂))) serves as a vehicle for compositional expressions, for idioms, and for regular fixed expressions like *het voordeel van de twijfel*. Hence, we may generalize the observations made so far, and propose three types of conceptual structures (see 12).

$$(12) \quad (\text{NP} (\text{art}) N_1 (\text{PP} \text{ prep}_3 ((\text{art}) N_2)))_X$$

a.	[OBJECT ₁] RELATION ₃ [OBJECT ₂]	<i>compositional</i>
b.	[OBJECT ₁ - OBJECT ₂] _X	<i>stored plus recoverable</i>
c.	[OBJECT] _X	<i>stored</i>

The differential representation of compositional (12a) and regular fixed expressions (12b) may account for the data we obtained in our experiment. Note that the nouns in the phrasal structure of (12) are represented as objects in (12a) and (12b). They are, therefore, available for further processing, e.g., as antecedents of anaphoric expressions, if the structures are integrated into the representation of the discourse. However, the objects in (12a) are 'free', whereas the objects in (12b) are fixed. This difference explains why anaphoric reference to one of the objects in (12a) is easier than to one of the objects in (12b). This is exactly what we found in our experiment. Assuming that the adjectival modification of regular fixed expressions like *het voordeel van de*

8. Compare similar Dutch expressions like *de nacht van de poezie* (the night of poetry), *de geschiedenis van de ruimtevaart* (the history of astronautics), *de rand van de afgrond* (the edge of the abyss), *de arrogantie van de macht* (the arrogance of the powers that be), *de maand van de filosofie* (the month of philosophy), *de vloek van pharao* (de curse of the pharao), and so on.

twijfel into *het voordeel van de lichte twijfel* renders them compositional, resolution of reference to *twijfel* takes more time and leads to more regressions in the former than in the latter case. It seems, then, that for regular fixed expressions the unified meaning, as indicated by the subscript X in (12b), supercedes the meaning of the constituting parts and is stored as such in the discourse representation, Retrieval of (one of) the objects is possible but time consuming.

An implication of this line of reasoning is that no such difference in processing should occur for unmodified and modified compositional expressions, like (12a). Our proposal thus predicts that anaphoric reference to *flat* should be equally easy for *de sleutels van de flat* (see 4) as for *de sleutels van de kleine flat* (*the keys to the small apartment*). This corollary of the model is the subject of our next experiment.

Unfortunately, the experiment has not shed light on the issue of how regular fixed expressions themselves are processed. The reading times of the expressions are slightly shorter for the modified versions (see Table 1), but the difference is not significant. The shorter reading times for the modified expressions may be attributable to the inserted adjective. Since the visual field in reading is skewed to the right (Rayner 1998), the noun *twijfel* may have been already read while the adjective *lichte* was fixated, which has resulted in shorter reading times of the noun and consequently of the complete expression. Because the adjective was absent in the unmodified condition, a true comparison is not possible, so the issue remains open.

Regarding the processing of modified expressions, two positions seem possible. The first is that if a modified expression is encountered, readers activate and store the parts of the expression in the discourse representation, but not its unified meaning, suggesting that the unified meaning hinges on the exact phrasing of the expression. The second position is that upon encountering the modified regular fixed expression, readers activate the unified meaning, insert the modifying adjective into its representation, thereby deactivating its unified meaning, and then store the result in the discourse representation. According to this position, the activation of the unified meaning hinges on the presence of all of its constituting parts. It seems that disentangling the two positions requires the determination of the exact moment during processing at which the unified meaning of a regular fixed expression becomes available (cf. the idiom key, see Tabossi & Zardon 1995). But we leave this issue to future research as well.

Appendix

Materials used in the experiment. The regular fixed expressions are in italics, the modifying adjective in parentheses. The analysis regions of the text are indicated by numbered slashes at the end of each region. The verification sentences are in italics below the texts.

- 1 Ik ben al langere tijd op zoek naar een interessante baan bij de overheid. Mijn voorkeur gaat uit naar een functie bij buitenlandse zaken.^{0/} Met ^{2/} (werkelijk) ^{3/} grote belangstelling ^{4/} las ik gisteren uw advertentie. ^{6/} In het bijzonder werd ^{7/} die ^{8/} gewekt door de eis Spaans te beheersen.^{9/} Ik spreek vloeiend Spaans en zou graag in het buitenland werken. Ik hoop dat ik voldoende gekwalificeerd ben voor de baan.^{11/}
Verification: De belangstelling werd gewekt door de eis Spaans te spreken.
- 2 Met het bestuur heb ik uw functioneren binnen onze firma besproken. Tijdens het overleg is ieders mening gehoord.^{0/} We willen u ^{1/} het voordeel van de ^{2/} (lichte) ^{3/} twijfel ^{4/} geven ^{5/} in deze bijzondere situatie.^{6/} Zoals u weet werd ^{7/} die ^{8/} opgeroepen door uw matige prestaties.^{9/} We hebben besloten u nog een kans te geven om aan uw gewijzigde werkomgeving te wennen. Tijdens de volgende functiebeoordeling bekijken we de zaak opnieuw.^{11/}
Verification: De twijfel werd opgeroepen door de matige prestaties.
- 3 Onlangs heeft u een huis gekocht in de Eggerlaan in Gouda. U heeft ons aangeschreven voor een vergunning voor een dakkapel.^{0/} Naar aanleiding van uw ^{2/} (uitgebreide) ^{3/} verzoek ^{4/} nemen wij contact met u op.^{6/} Voorlopig kunnen wij ^{7/} het ^{8/} helaas niet in behandeling nemen.^{9/} We ondervinden namelijk veel last van een interne verhuizing. Binnen afzienbare tijd zullen wij u een reactie geven.^{11/}
Verification: Het verzoek kan voorlopig niet in behandeling worden genomen.
- 4 De boekenbeurs vindt dit jaar voor de vijftiende keer plaats. Ieder jaar nodigt de organisatie een bekende schrijver uit als gastspreker.^{0/} Het is mij een ^{2/} (bijzonder) ^{3/} groot genoegen^{4/} u dit jaar te mogen verwelkomen ^{6/}. Deze keer is ^{7/} het ^{8/} met name groot ^{9/}, omdat ik het evenement voor de laatste keer organiseer.^{10/} Ik vind het prettig om met een indrukwekkend programma af te sluiten. Met uw komst als gastspreker ben ik daarvan verzekerd.^{11/}
Verification: Het genoegen is met name deze keer groot.
- 5 Ik zou graag willen dat uw zoon zich fatsoenlijk gedraagt. Dat geldt zowel in de klas, in de kantine als op het schoolplein.^{0/} Hij werkt op mijn ^{2/} (gevoelige) ^{3/} zenuwen ^{4/} met zijn vreselijk asociale gewoontes.^{6/} De laatste tijd worden ^{7/} ze ^{8/} zo geprikkeld dat ik moeite heb me in te houden ^{9/}. Ik wil dat hij zijn gedrag binnen een week verandert. Anders neem ik contact op met de rector.^{11/}
Verification: De zenuwen worden de laatste tijd erg geprikkeld.

- 6 Ik heb u vorige week om advies gevraagd over de bevalling van mijn konijn. Binnen tien minuten kreeg ik al een e-mail van u terug.^{0/} (Bijzonder) ^{3/} *Hartelijk dank* ^{4/} voor de erg snelle reactie.^{6/} Zoals u begrijpt is ^{7/} deze ^{8/} erg groot ^{9/}, omdat ik niet wist hoe te handelen.^{10/} Zonder uw advies had mijn konijn de bevalling misschien niet overleefd. Ze maakt het nu goed en zal snel hersteld zijn.^{11/}
Verification: De dank was erg groot.
- 7 Ik heb een nogal vervelende mededeling voor u. Het is voor mij helaas niet mogelijk om vanavond op de vereniging aanwezig te zijn.^{0/} Ik heb vandaag mijn ^{2/} (beste) ^{3/} *dag* ^{4/} *niet* ^{5/} en laat de vergadering schieten.^{6/} Volgens mij was ^{7/} deze ^{8/} al slecht toen ik opstond.^{9/} Ik werd wakker met hoofdpijn en verstuikte mijn enkel bij het opstaan. Zodra ik me beter voel neem ik contact met u op.^{11/}
Verification: De dag was al slecht bij het opstaan.
- 8 U heeft via de website privé-foto's verspreid van een van onze directiemedewerkers. Daarom mag u de komende tijd niets op de website plaatsen.^{0/} Hopelijk ^{1/} heeft u uw ^{2/} (pijnlijke) ^{3/} *lesje* ^{4/} *geleerd* ^{5/} en verandert uw gedrag ^{6/}. Wij raden u aan ^{7/} dat ^{8/} goed tot u te nemen.^{9/} De directie is zeer geschokt door dit vervelende voorval. Het is op zijn plaats uw excuses aan te bieden aan de betreffende persoon.^{11/}
Verification: Het lesje moet goed onthouden worden.
- 9 Ik weet dat de studie Econometrie niet makkelijk is. Veel mensen haken al na enkele maanden af omdat de stof te moeilijk blijkt.^{0/} Toch moet je ^{1/} de ^{2/} (zware) ^{3/} *rit* ^{4/} *uitzitten* ^{5/} en met volharding doorwerken.^{6/} Je zult zien dat ^{7/} deze ^{8/} achteraf de moeite waard is.^{9/} De studie biedt je veel kans op een goede baan. En die zijn tegenwoordig erg moeilijk te krijgen.^{11/}
Verification: De rit zal achteraf de moeite waard zijn.
- 10 Het dragen van een helm op de bouw is verplicht. Desondanks dragen velen van jullie nog steeds geen helm.^{0/} Jullie betreden ^{1/} met ^{2/} (ernstig) ^{3/} *gevaar* ^{4/} *voor eigen leven* ^{5/} je werkplek.^{6/} Zeer waarschijnlijk schatten jullie ^{7/} dit ^{8/} als heel klein in.^{9/} Toch overlijden jaarlijks tientallen mensen als gevolg van bedrijfsongelukken. In veel van de gevallen droegen de slachtoffers geen hoofdbescherming.^{11/}
Verification: Het gevaar wordt waarschijnlijk als zeer klein ingeschat.

- 11 Ik richt me tot u namens de directie. Mede door de huidige economische recessie gaat het niet goed met ons bedrijf en dat baart ons zorgen.^{0/} Het is *van* ^{2/} (bijzonder) ^{3/} *groot belang* ^{4/} dat iedereen veel werk verzet.^{6/} Op dit moment is ^{7/} dit ^{8/} vooral groot ^{9/} omdat de maandcijfers zeer ongunstig zijn.^{10/} Ik vraag u dan ook om een extra inspanning te leveren. Alleen dan is het mogelijk om zonder ontslagen deze periode door te komen.^{11/}
Verification: Het belang is vooral op dit moment groot.
- 12 Volgende week zaterdag is het open dag op de muziekschool in Vlissingen. Net als andere jaren hebben we weer veel hulp van vrijwilligers nodig.^{0/} *Langs deze* ^{2/} (originele) ^{3/} *weg* ^{4/} willen wij u vragen om hulp.^{6/} Wij hopen dat ^{7/} hij ^{8/} tot veel inschrijvingen leidt.^{9/} We verwachten veel bezoekers, dus uw hulp is hard nodig. Zoals gewoonlijk wordt iedere vrijwilliger beloond met een leuk presentje.^{11/}
Verification: De weg leidt hopelijk tot veel inschrijvingen.
- 13 Als bedrijfsarts heb ik nu drie gesprekken met u gevoerd. Ik vind dat u uw problemen niet langer voor u moet houden. Dat lost niets op.^{0/} *Gooi het maar in de* ^{2/} (gehele) ^{3/} *groep* ^{4/} bij de volgende gelegenheid.^{6/} U moet ervoor zorgen dat ^{7/} deze ^{8/} ook op de hoogte is.^{9/} Het is beter dat uw collega's weten wat er aan de hand is. Dan kunnen zij rekening houden met uw situatie.^{11/}
Verification: De groep moet ook op de hoogte zijn.
- 14 U heeft mij een grote dienst bewezen. Nu ik een tuinhek heb mogen plaatsen is het een stuk veiliger geworden, met name voor onze kinderen.^{0/} *Uit de grond van mijn* ^{2/} (bezwaarde) ^{3/} *hart* ^{4/} dank ik voor uw begrip.^{6/} Al geruime tijd was ^{7/} dit ^{8/} behoorlijk onder druk komen te staan gezien uw afkeer van hekken.^{9/} Ik weet dat u liever een vrij uitzicht had gehad. Ik zal u zeker nog belonen voor uw toegeeflijkheid.^{11/}
Verification: Het hart was onder druk komen te staan.
- 15 De rapportcijfers van uw zoon zijn zorgwekkend. Hij staat voor vijf vakken onvoldoende en lijkt daar niet ongerust over te zijn.^{0/} Uw zoon ^{1/} heeft *een goed stel* ^{2/} (werkende) ^{3/} *hersens* ^{4/} en volgt tevens alle lessen ^{6/}. Op dit moment worden ^{7/} ze ^{8/} echter voor andere doeleinden ingezet dan de studie.^{9/} Als hij op deze manier doorgaat zal hij zeker blijven zitten. Ik raad u aan om hierover met hem te praten.^{11/}
Verification: De hersenen worden voor andere doeleinden gebruikt dan de studie.

- 16 In februari ben ik afgestudeerd in Vrijetijdwetenschappen aan de UvT. Sindsdien ben ik op zoek naar een functie die bij me past.^{0/} *Naar aanleiding van uw* ^{2/} (opvallende) ^{3/} *advertentie* ^{4/} schrijf ik u aan.^{6/} Op de volle pagina was ^{7/} deze ^{8/} zo opvallend dat ik er niet omheen kon.^{9/} U bent op zoek naar een beleidsmedewerker voor uw pretpark. Het lijkt me leuk om bij een pretpark te werken.^{11/}
Verification: De advertentie was opvallend op de volle pagina.
- 17 Voor u ligt de conceptversie van het vernieuwde clubboekje. Zoals ieder jaar hebben we er weer veel tijd en energie in gestoken.^{0/} *Mocht u* ^{2/} (belangrijke) ^{3/} *op- of aanmerkingen* ^{4/} *hebben* ^{5/}, schrijf dan de commissie aan ^{6/}. Wij verzoeken u ^{7/} deze ^{8/} meteen door te geven ^{9/}, dan kan het boekje snel naar de drukkerij.^{10/} Het zal waarschijnlijk begin juni gedrukt zijn. U kunt dan voor vijf euro een exemplaar kopen in de kantine.^{11/}
Verification: De op- of aanmerkingen moeten meteen doorgegeven worden.
- 18 Je bent met lof geslaagd voor de studie Scheikundige Technologie aan de TU in Eindhoven. Na jaren studeren ben je nu klaar voor de arbeidsmarkt.^{0/} Ik wens je ^{2/} (uitermate)^{3/} *veel succes* ^{4/} met het vinden van werk ^{6/}. Vooral nu kun je ^{7/} dat ^{8/} goed gebruiken ^{9/}, want er zijn niet veel vacatures.^{10/} Het gaat slecht met de economie en daar heeft iedereen last van. Maar gezien je uitzonderlijk capaciteiten zul je waarschijnlijk toch snel een plekje verwerven.^{11/}
Verification: Vooral nu kun je succes goed gebruiken.

References

- Bolinger, D. 1976. Meaning and memory. *Forum Linguisticum* 1: 1–14.
- Chafe, W. 1968. Idiomaticity as an anomaly in the Chomskyan Paradigm. *Foundations of Language* 4: 109–127.
- Cozijn, R. 2006. Het gebruik van oogbewegingen in leesonderzoek. *Tijdschrift voor Taalbeheersing* 28(3): 220–232.
- Dabrowska, E. 2004. Rules or schemas? Evidence from Polish. *Language and Cognitive Processes* 19: 225–271.
- De Coster, M. 2002. *Woordenboek van Populaire Uitdrukkingen, Clichés, Kreten en Slogans* Den Haag: SDU Uitgevers.
- Deane, P.D. 1991. Limits to attention: A cognitive theory of island phenomena. *Cognitive Linguistics* 1:1–63.
- Erman, B. & Warren, B. 2000. The idiom principle and the open choice principle. *Text* 20: 29–62.
- Everaert, M., van der Linden, E.J., Schenk, A. & Schreuder, R. (eds). 1995. *Idioms: Structural and Psychological Perspectives*. Hillsdale NJ: Lawrence Erlbaum Associates.

- Frazier, L. & Rayner, R. 1982. Making and correcting errors during sentence comprehension: Eye movements in the analysis of structurally ambiguous sentences. *Cognitive Psychology* 14: 178–210.
- Geeraerts, D. 1995. Specialization and reinterpretation in idioms. In *Idioms: Structural and Psychological Perspectives*, M. Everaert, E.J. van der Linden, A. Schenk & R. Schreuder (eds), 57–74. Hillsdale NJ: Lawrence Erlbaum Associates.
- Gibbs, R.W. 1995. Idiomaticity and human cognition. In *Idioms: Structural and Psychological Perspectives*, M. Everaert, E.J. van der Linden, A. Schenk & R. Schreuder (eds), 97–116. Hillsdale NJ: Lawrence Erlbaum Associates.
- Goldberg, A.E. 1995. *Constructions: A Construction Grammar Approach to Argument structure*. Chicago IL: University of Chicago Press.
- Jackendoff, R.S. 1997. *The Architecture of the Language Faculty*. Cambridge MA: The MIT Press.
- Jackendoff, R.S. 2002. *Foundations of Language: Brain, Meaning, Grammar, Evolution*. Oxford: OUP.
- Jackendoff, R.S. & Pinker, S. 2005. The nature of the language faculty and its implications for evolution of language (Reply to Fitch, Hauser, and Chomsky). *Cognition* 97(2): 211–225.
- Kuiper, K. 1996. *Smooth Talkers: The Linguistic Performance of Auctioneers and Sportscasters*. Mahwah NJ: Lawrence Erlbaum Associates.
- Kuiper, K., van Egmond, M.E., Kempen, G. & Sprenger, S.A. 2007. Slipping on superlemmas: Multiword lexical items in speech production. *The Mental Lexicon* 2: 313–357.
- Langacker, R.W. 1988. A usage-based model. In *Topics in Cognitive Linguistics*, B. Rudzka-Ostyn (ed.), 127–164. Amsterdam: John Benjamins.
- Langacker, R.W. 2000. A dynamic usage-based model. In *Usage Based Models of Language*, M. Barlow & S. Kemmer (eds), 1–64. Cambridge: CUP.
- Meščuk, I. 1995. Phrasemes in language and phraseology in linguistics. In *Idioms: Structural and Psychological Perspectives*, M. Everaert, E.J. van der Linden, A. Schenk & R. Schreuder (eds), 167–232. Hillsdale NJ: Lawrence Erlbaum Associates.
- Murray, W.S. 2000. Sentence processing: Issues and measures. In *Reading as a Perceptual Process*, A. Kennedy, R. Radach, D. Heller & J. Pynte (eds), 649–666. Oxford: Elsevier.
- Nicolas, T. 1995. Semantics of idiom modification. In *Idioms: Structural and Psychological Perspectives*, M. Everaert, E.J. van der Linden, A. Schenk & R. Schreuder (eds), 233–252. Hillsdale NJ: Lawrence Erlbaum Associates.
- Pinker, S. 1994. *The Language Instinct*. New York NY: W. Morrow and Co.
- Pinker, S. & Jackendoff, R.S. 2005. The faculty of language: What's special about it? *Cognition* 95(2): 201–236.
- Pollatsek, A. & Well, A.D. 1995. On the use of counterbalanced designs in cognitive research: A suggestion for a better and more powerful analysis. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 21: 785–794.
- Rayner, K. 1998. Eye movements in reading and information processing: 20 Years of research. *Psychological Bulletin* 12: 372–422.
- Schilperoord, J. 1996. *It's About Time: Temporal Aspects of Cognitive Processes in Text Production*. Amsterdam: Rodopi.
- Schilperoord, J. & Peeters, J. 1998. Reconstructie en completering van gefixeerde en niet-gefixeerde expressies. Ms, Tilburg University.
- Sprenger, S.A. 2003. Fixed expressions and the production of idioms. Ph.D. dissertation, Radboud University, Nijmegen.

- Tabossi, P. & Zardon, F. 1995. The activation of idiomatic meaning. In *Idioms: Structural and Psychological Perspectives*, M. Everaert, E.J. van der Linden, P.J. Schellens, A. Schenk & R. Schreuder (eds), 273–282. Hillsdale NJ: Lawrence Erlbaum Associates.
- Van de Voort, E.C. & Vonk, W. 1995. You don't die immediately when you kick an empty bucket: A processing view on semantic and syntactic characteristics of idioms. In *Idioms: Structural and Psychological Perspectives*, M. Everaert, E.J. van der Linden, A. Schenk & R. Schreuder (eds), 283–300. Hillsdale NJ: Lawrence Erlbaum Associates.
- Van Gestel, F. 1995. En bloc insertion. In *Idioms. Structural and Psychological Perspectives*, M. Everaert, E.J. van der Linden, A. Schenk & R. Schreuder (eds), 75–96. Hillsdale NJ: Lawrence Erlbaum Associates.
- Vonk, W. & Cozijn, R. 2003. On the treatment of saccades and regressions in eye movement measures of reading time. In *The Mind's Eye: Cognitive and Applied Aspects of Eye Movement Research*, J. Hyönä, R. Radach & H. Deubel (eds), 291–312. Oxford: Elsevier.
- Wray, A. & Perkins, M. 2000. The functions of formulaic language: An integrated model. *Language and Communication* 20: 1–28.

Clitic doubling in Spanish

Agreement of the third kind

Jan Schroten
Utrecht University

The syntactic behavior of nouns is in part determined by their meaning; for example, nouns which refer to human beings usually have special syntactic properties. The basic distinction between human and non-human nouns is a psychological reality reflected in language specific properties of many languages. The question is how this difference is reflected and why the distinction surfaces the way it does. This is the basic issue that will be discussed in this contribution, by taking a look at a much discussed property of Spanish. In Spanish, the definite human direct objects can be doubled by a clitic pronoun in many dialects, and must be doubled if it is a strong, human pronoun. Definite human direct objects must be introduced by the preposition *a* in all variants, and can be doubled by a clitic in many variants. The question what makes human direct objects so special and which other properties can be found has been studied very much in different regions and dialects or variants. In this paper, attention will be given to the other side of the question: What blocks clitic doubling if the direct object is non-human? Since clitic doubling is a kind of agreement, the question is why and how clitic doubling is blocked with non-human objects.

1. Introduction

In Spanish, clitic doubling is a common phenomenon that has been discussed in different contexts. In order to see what clitic doubling is, let us take a look at the examples (1)–(3):

- (1) El niño saluda a la profesora
(‘The child greets to the female.teacher’)
- (2) El niño la saluda
(‘The child her greets’)
- (3) El niño la saluda a la profesora
(‘The child her greets to the female.teacher’)

The verb form *saluda* (‘greets’) of transitive *saludar* (‘to greet’) combines with a direct object (DO) in these three examples. The DO in (1) is *a la profesora* (‘to the female.

teacher'), a Determiner Phrase (DP), which is composed of the feminine singular definite article *la* ('the') and its complement, the Noun *profesora* ('female.teacher') and which is introduced by the preposition-like Case marker *a* ('to'). This full DO follows the verb. The DO in (2) is *la* ('her'), which is a clitic pronoun or clitic for short, unstressed and located at the left edge of the tensed verb form, as a kind of prefix. In (3), we find clitic doubling: the DO is expressed twice, by the Case-marked DP *a la profesora* ('to the female.teacher') and by the DO clitic pronoun *la* ('her'). Thus, the clitic *la* ('her') doubles the full DO, the Case-marked DP *a la profesora* ('to the female.teacher') or the clitic can be said to be doubled, and specified in interpretation, by *a la profesora* ('to the female.teacher').

The full, non-clitic, DO in (1) and (3) is introduced by the preposition-like Case marker *a* ('to'), which is obligatory with full definite human DOs in Spanish. Omission of the Case marker gives unacceptable and ungrammatical results, as shown in (1') and (3'):

- (1') *El niño saluda la profesora
('The child greets the female.teacher')
- (3') *El niño la saluda la profesora
('The child her greets the female.teacher')

Clitic doubling in Spanish is clearly related to the animacy, that is, reference to one or more human beings, of the DO. Non-animate or non-human DOs have other referents, such as things or objects, and processes and other abstract referents and these DOs are not introduced by the preposition-like accusative Case-marker *a* and they do not permit clitic doubling, as is shown in (4):

- (4) *Case marking and clitic doubling of the full non-human DO*
- (a) El niño vio (*a) la mesa
('The child saw (to) the table')
- (b) El niño (*la) vio la mesa
('The child (her) saw the table<FEM>')

Clitic doubling has been investigated as an exceptional property of Spanish which needs explanation, by linking it to another exceptional property, the preposition-like Case marker *a* ('to') that introduces definite human full DOs. Historically related languages such as French, Italian and Portuguese are similar in that the full DO follows the verb, and that the clitic pronoun is like an affixal element of the verb, but they are different in that clitic doubling is not used.

In this paper, we will take a look at the other side of clitic doubling. There is no doubt that clitic doubling is an instance of agreement, which holds between the DO clitic and its double, the full DO, which is a DP. The other side of clitic doubling, then, is the question why clitic doubling is not always possible with full DOs. In other words,

the question is what blocks DO clitic doubling with non-human DOs, as already shown in (4b).

In order to be able to discuss this question, the following assumption will be made:

(5) *The status of clitic pronouns in Spanish*

Spanish clitic pronouns are affixes that are part of the verb

In the minimalist framework that will be adopted in this study (Chomsky (1995, 2000, 2001)), the obligatory character of subject-verb agreement is analyzed as the effect of the uninterpretability of the person-number features which are visible in the suffix of the tensed verb. If the subject has the same values for person and number as the affix of the tensed verb, it erases these features in the suffix. If clitic pronouns are affixes, they can be interpreted as having the same properties, which are given in (6):

(6) *Spanish clitic pronouns in the minimalist framework*

(i) The number-gender features of the clitic pronoun are uninterpretable in its location as an affix of the verb;

(ii) Therefore, these visible features of the clitic pronoun must be erased.

The intuitive basis of the analysis of clitic doubling in this paper is that it is agreement with exceptional properties, which are due to the fact that the gender of the Noun, which is masculine or feminine in Spanish, has no meaning. For example, the lexical feature feminine (which will be represented as ⟨fem⟩) of the Noun *mesa* ('table') has effects on the company it keeps, such as the definite article *la* and the adjective *bonita* in the phrase *la mesa bonita* ('the table nice'), where *la* ('the'), *mesa* ('table') and *bonita* ('nice') manifest their feminine character by the form they have. Within this phrase, the various manifestations of singular number and feminine gender can be assumed to be erased, except for the highest ones in the phrase. The remaining, highest, gender feature, however, is uninterpretable if the Noun has a non-human referent, such as a thing. The things denoted by feminine Nouns have no semantic or pragmatic property in common. However, if the Noun denotes a human being, that is, if it is a human or animate Noun, the gender feature is nearly always interpretable: feminine human Nouns refer to persons of the female sex, and masculine human Nouns to persons of the male sex or to persons of both sexes, since the masculine gender is the default, weak gender.

The effects of the interpretability and uninterpretability of the gender feature of Nouns, pronouns and clitics on agreement will be sketched in the next sections.

In the second section, the clitics and the strong, full or tonic pronouns and their behavior will be sketched, and how to treat and represent their interpretation. In the third section, types of agreement and agreeing features in DO clitic doubling and subject-verb agreement will be treated. A sketch will be given of how to account for the impossibility of doubling the non-human DO. In the fourth section, indirect object

clitic doubling, IO clitic doubling, will be taken into consideration, and compared with subject-verb agreement and with DO clitic doubling. In the fifth section, some observations will be made on the different behavior with respect to clitic doubling that we find in historically related languages, French and Italian. In the final section, a short résumé will be given.

2. Spanish clitic pronouns and strong pronouns and their interpretation

2.1 Spanish clitic pronouns: Their status and their interpretation

The Spanish clitic pronouns are unstressed and they cannot be stressed contrastively. Phonologically, they are part of the verb. A plausible assumption, then, is that clitics are affixes related to the (direct or indirect) object, in the same way as the person-number suffixes of the tensed verb are related to the subject. The affixes, which are part of the verb, cannot be interpreted *in situ* in the minimalist framework. They should match the values of the relevant features in the related argument position.

There are two basic generative pre-minimalist approaches to the clitics. One approach is to generate the clitic in the DO position following the verb and to move it to the preverbal position of tensed verbs, leaving a trace in the original position where interpretation is done. The other approach is to generate it in clitic position, that is, in a position generated by a functional projection located close to the verb, and to assume that there is a zero pronoun in the postverbal DO position, where interpretation is done by combining it with the clitic antecedent. The clitic doubling phenomenon found in Spanish favors the second approach, and has been used as an argument against the first approach (cf. Jaeggli 1993; Mendikoetxea 1993; Suñer 1993; Suñer 1999).

Subject-verb agreement is the matching of the number and the person values of the subject and of the suffix of the tensed verb. In a similar fashion, DO clitic doubling can be taken to be the matching of number and gender of the clitic and the number and gender of the full DO. In short, DO clitic doubling can be analyzed as number and gender agreement of the DO-clitic and the full DO.


The regular third person direct object clitic of Spanish has visible gender and number:

(7) *Forms and features of the Spanish third person DO clitic*

- (i) *lo*: [MASC]&[SG], that is: masculine and singular;
- (ii) *la*: [FEM]&[SG], that is, feminine and singular;
- (iii) *los*: [MASC]&[PL], that is, masculine and plural;
- (iv) *las*: [FEM]&[PL], that is, feminine and plural.

If the clitic is generated as an affix of the verb, uninterpretable in that position, its interpretation requires the assumption that the clitic, or its visible features, is copied into the usual DO-position to the right of the verb, as sketched in diagram (8):

(8) *Tree-like diagram of la vio ('her (s)he.saw')*



In order for the features of the DO clitic to be interpretable, they must be copied into the DO position, where I assume that they take the Determiner position D of the direct object DP, as sketched in (8). Consider now what it means for the clitics to be interpreted in the D-position of the full DO. Both the number feature and the gender feature can be taken to be some kind of operator which ranges over the Noun which heads the complement NP of D in DO position. The copy of number and gender in the DO Determiner makes it possible to erase these features in the clitic position where they came from, and where they are uninterpretable.

This copying procedure is sketched in (9):

(9) *The DO-clitic and its copy in DO position*

(i) First step:

Copying the values of number and gender of the clitic in DO position

[V: [clitic: [α_{FEM}] & [β_{PL}]]+V] – [DP:[D:[+D]][NP:[N [+N]]]]

⇒ [V:[clitic: [α_{FEM}] & [β_{PL}]]+V] – [DP:[D:[+D, [α_{FEM}] & [β_{PL}]]][NP:[N [+N]]]]

(ii) Second step: Erasing the uninterpretable features of the clitic in V

[V:[clitic: [α_{FEM}] & [β_{PL}]]+V] – [DP:[D:[+D, [α_{FEM}] & [β_{PL}]]][NP:[N [+N]]]]

⇒ [V:[clitic]+V] – [DP:[D:[+D, [α_{FEM}] & [β_{PL}]]][NP:[N [+N]]]]

NOTE: α = + or –; β = + or –

Now note that the DP in DO position still has no meaning, that is, has no identified referent, since the complement of the filled D has no content, neither phonological nor semantic. Its referent is found in discourse or by pragmatic context, by a suitable antecedent or by a suitable object which can be pointed at, and which can be named by a suitable lexical item. The antecedent or the object is suitable if it is, or if it is associated with, a lexical item having the same values for gender and number.

I will use the question mark to indicate that the phrase has been interpreted, in some sense, tautologically. It will have a real referent if a matching antecedent DP with matching values for number and gender is found or an object which can be referred to by such a matching DP.

Thus, the clitic *la* ('her') in the phrase: *La vi* ('Her I.saw'), whose structure has been given in (8), is interpreted in the following way. Its features are copied into the D-position of the null postverbal direct object DP and its null N complement is interpreted as waiting for a suitable antecedent or pragmatically defined referent, and this interpretation is represented by a question mark, as in (10):

- (10) *Interpretation of the clitic la in the postverbal direct object DP*
 [D:([-PL](x)&[+FEM](x)) [N: ⟨x = ?⟩]

The lack of meaning of the copied clitic is due to the lack of content of the NP. This can be hypothesized to be the reason why the clitic cannot get emphatic stress: there is no interpretable feature that can be given emphasis. Neither the number feature nor the gender feature can be stressed to give prominence to the referent, which must be found in the Noun.

2.2 Spanish strong pronouns: Their status and their interpretation

Now consider the forms and features of the third person strong pronoun:

- (11) *Forms and features of the Spanish third person DO strong pronoun*
- (i) *él*: [MASC]&[SG], that is: masculine and singular;
 - (ii) *ella*: [FEM]&[SG], that is, feminine and singular;
 - (iii) *ellos*: [MASC]&[PL], that is, masculine and plural;
 - (iv) *ellas*: [FEM]&[PL], that is, feminine and plural.

Strong DO pronouns have the same features as the DO clitics, but the forms are different and they bear stress, which is why they are strong. Their usual position is that of the full direct object, the postverbal position.

The Spanish strong DO pronoun refers obligatorily to a human being (cf. Fernández Soriano 1999). The accusative, DO pronoun, then, has the interpretable nominal feature [+human]. Therefore, it can be assumed that the strong DO pronoun is not a Determiner, but a Noun. If this pronoun is a Noun, it needs a Determiner, an operator which takes the semantic contents of the strong pronoun in its range. An economic way of providing the Determiner with some operator is to use a matching third person clitic pronoun, that is, by clitic doubling the DO. And this is what we find in all variants of Spanish: the direct object strong pronoun with human reference requires a clitic, which agrees in gender and number. Of course, it needs also the prepositional Case marker *a* ('to'), which gives the paradigm shown in (12):

- (12) *Using the strong human direct object pronoun and its consequences*
- (a) **Vi ella* ('I.saw her')

- (b) *Vi a ella ('I.saw to her')
- (c) La vi a ella ('Her I.saw to her')

Thus, the determiner D has a number operator and a gender operator; the number operator is as a quantifier, and can be taken to be interpretable as such. The gender feature by itself, however, cannot be used to restrict the interpretation of things that are referred to by a Noun having masculine or feminine gender. There simply is no way of giving interpretation to the gender of nouns with non-human interpretation. The only way of making the gender feature interpretable is to use it as an operator which ranges over human nouns, interpreting masculine human nouns as referring to male persons and female nouns as referring to female persons. As usual, the masculine gender is the default gender and masculine human Nouns can refer to persons of the male and the female sex (cf. Schroten 2001a, 2001b).

In subject position, the strong pronoun, which has the same forms and interpretation as the strong DO pronoun (see (11)) can only be used if it has human reference. Of course, it must agree with the person-number affix of the tensed verb, whose values can be assumed to be copied into the Determiner position of the subject, the same way as sketched for the clitic in (10). The strong pronoun can be omitted, since Spanish is a null subject language. The null subject is a DP, and the values of the affix are copied into the D-position, taking the null Noun as its complement where it is interpreted as requiring a suitable antecedent or pragmatically identified object. This is sketched in (13):

(13) *The subject pronoun, its interpretation and the interpretation of the null subject*

- (a) Ella está<3sg> ahí ('She is there')
 Interpretation of the subject DP:
 [DP: [D: <3sg>] [N: ella <+human> & <FEM>]] => <the woman>
- (b) Está <3sg> ahí ('Is there')
 Interpretation of the null subject:
 [DP: [D: <3sg>] [N: <x = ?>]]

To resume, this is how obligatory clitic doubling of the strong pronominal direct object, found in all variants of Spanish, can be treated as an extension of its null-subject properties. The person-number values of the affix of the tensed forms are copied into subject position, in Determiner position with a null NP complement. Again, the referent of the null subject is impossible to identify if the affix has third person value and singular or plural number. However, reference can be established in discourse, by identifying an appropriate antecedent or by pointing to a suitable object. As with DO clitics, the third person null subject referent is a question mark in the first step of interpretation, to indicate that the referent can only be identified by finding a suitable antecedent or a suitable object in the appropriate discourse context.

If a third person strong pronoun is used as a subject, it can only be used if it has human reference. Again, it can be hypothesized that the lexical pronoun is a noun needing a Determiner, which is supplied by the person-number affix of the verb. If the

strong pronoun in subject position is just a Noun, the Determiner operator number has no suitable range. The most basic semantic-referential value that a Noun can have is the interpretable feature [+human].

Given the similarity of null subjects of tensed verbs and the direct object clitic, it can be concluded that they are both weak pronouns (cf. Cardinaletti, A.; Starke, M. (1999)) lacking semantic features which make it possible to identify the referent without the intervention of a suitable antecedent in discourse or a suitable object. The associated lexical strong pronoun is a Noun with the interpretable feature [+human].

3. Types of agreement and agreeing features

The most common type of agreement is subject-verb agreement, which is the use of the same values for person and number in the subject and the affix of the tensed verb. Considering third person subjects, agreement is matching of number: singular subjects match [3sg] – third person singular – affixes of the tensed verb forms, and plural subjects match [3pl] – third person plural – affixes of the tensed verb forms, as in:

(14) *Third person subject verb agreement*

- (a) El pájaro cantó / *cantaron ('The bird sang<SG> / sang <PL>')
- (b) Los pájaros cantaron / *cantó ('The birds sang <PL> / sang <SG>').

In DO clitic doubling constructions, the clitic agrees in number and gender with the full DO, as in:

(15) *Third person DO clitic doubling agreement*

- (a) La / *Lo vi a la chica ('Her/Him I.saw to the girl')
- (b) Lo / *La vi al chico ('Him/Her I.saw to the boy')
- (c) Las /*Los vi a las chicas ('Them<FEM>/Them<MASC> I.saw to the girls')
- (d) Los /*Las vi a los chicos ('Them<MASC>/Them<FEM> I.saw to the boys')

Aspects of syntactic and semantic issues with respect to clitic doubling of [+human] direct objects have been discussed in detail by Gutiérrez Rexach (2001) and by Ormazabal and Romero (2007).

Now consider the basic difference between subject-verb agreement and DO doubling. The agreeing features in subject-verb agreement, which involves number and person, are not the same features as in direct object clitic doubling, which involves number and gender. Subject-verb agreement is possible and obligatory with any kind of lexical subject, independent of the [+human] vs. [–human] semantic content of the subject, since the affix on the verb has no visible gender feature. In DO clitic doubling, the clitic and the DO agree in person and gender when agreement is found. But [–human] direct objects cannot be doubled by the direct object clitic.

The difference, then, is that the person-number features of the affix of the tensed verb can be erased by the person-number feature values of the subject, which are interpreted and have the strength to erase the values on the affix.

Now, if the direct object tries to erase the gender-number feature of the clitic, its gender feature is uninterpretable. This uninterpretability of the gender feature will be represented as value zero. Note that the gender value ⟨fem⟩ or ⟨masc⟩ will remain visible in the course of derivation, but its interpretation adds nothing to its semantics. The hypothesis that emerges is that zero-interpretation makes the direct object DP lose vitality. It is unable to erase the gender feature of the direct object clitic. This hypothesis on the interpretation of the gender feature of [–human] Nouns is sketched in (16):

(16) *Interpreting the gender of non-[HUMAN] Nouns*

The gender feature [MASC] or [FEM] of the [–human] Noun is interpreted as zero, represented as [Ø].

What must be added is the assumption that zero-interpreted features cannot erase matching features: they are weak. This hypothesis is stated in (17):

(17) *The strength of zero-interpreted features*

Zero-interpreted features are weak, and unable to erase a matching feature.

The derivation and final crash of the clitic doubled non-human Noun in: **La vi la mesa* (Her I-saw the ⟨SG⟩&⟨FEM⟩ table ⟨SG⟩&⟨FEM⟩) is sketched in (16):

(18) *The crash of a clitic doubled [–human] full direct object*

(i) Step 1:

Erase matching features in the DP, keeping the highest ones:

$la_{\langle FEM \rangle} \& \langle SG \rangle + mesa_{\langle FEM \rangle} \& \langle SG \rangle$ ('the + table')

⇒ $[DP_{\langle FEM \rangle} \& \langle SG \rangle: la + mesa]$

(ii) Step 2: Interpret the features of DP:

$[DP_{\langle FEM \rangle} \& \langle SG \rangle: [[D: [la]] + [NP: [N: [mesa]]]]]$

⇒ $[DP_{\langle FEM = \emptyset \rangle} \& \langle SG = \text{Quantifier (1)} \rangle: la + mesa]$

(iii) Step 3: Combine the Verb, including the DO-clitic and DO DP:

⇒ $[V: la_{\langle FEM \rangle} \& \langle SG \rangle + V\text{-stem} + \text{affix}] +$
 $+ [DP_{\langle FEM = \emptyset \rangle} \& \langle SG = \text{Quantifier (1)} \rangle: la + mesa]$

(iv) Step 4:

Erase the features of the clitic which match the active features of DP

$[V: la_{\langle FEM \rangle} \& \langle SG \rangle + V\text{-stem} + \text{affix}] +$
 $+ [DP_{\langle FEM = \emptyset \rangle} \& \langle SG = \text{Quantifier (1)} \rangle: la + mesa]$

⇒ $[V: la_{\langle FEM \rangle} + V\text{-stem} + \text{affix}] +$
 $+ [DP_{\langle FEM = \emptyset \rangle} \& \langle SG = \text{Quantifier (1)} \rangle: la + mesa]$

(v) Step 5:

The ⟨FEM⟩ gender feature of the clitic *la* is uninterpretable:

The derivation crashes: **la_{\langle FEM \rangle} vi la mesa*

Of course, this approach has to be elaborated, since the gender feature of the direct object DP, although it is weak, must be visible in discourse, taking an appropriate antecedent.

The main reason to choose this approach is that DPs containing a zero interpreted gender feature have other peculiar properties: they cannot be stressed and are doomed to inactivity. Also in other languages with uninterpretable gender of [–human] Nouns, such as Dutch, these weak pronouns cannot be emphasized (see Cardinaletti & Starke 1999; Corver and Delfitto 1999).

A question which cannot be discussed now is whether direct object clitic-doubling still is active when there is no visible clitic by selecting a null clitic with erasable features only, such as number, or whether direct object agreement is optional. The route that has been chosen would lead us to expect that the null clitic option is the most promising one. However, I am unable to give support to this intuition.

4. Indirect object clitic doubling in Spanish

An issue that we have not taken into account until now is indirect object clitic-doubling, as we find in (19a) and (19b):

- (19) a. Les gustan los tangos a los argentinos
(‘Them give.pleasure the tangos to the Argentines’)
b. Les faltan algunas páginas a estos libros
(‘Them lack some pages to these books’)

In both examples, the omission or zeroing of the indirect object clitic is awkward and can be qualified as ungrammatical (cf. Jaeggli 1993):

- (19') a. *Gustan los tangos a los argentinos
(‘Give.pleasure the tangos to the Argentines’)
b. *Faltan algunas páginas a estos libros
(‘Lack some pages to these books’)

An interesting property of third person indirect object clitics is that they express person and number, but that gender is not expressed. In this sense, it is like subject-verb agreement, and there is no reason, as there is with direct object clitics, not to have visible agreement. And in fact, indirect object clitic doubling is common, and in many cases obligatory, even if the doubled indirect object has [–human] reference, as has been shown in (19b), (19'b).

5. On clitic doubling in other Romance languages

Given the approach to clitic doubling in Spanish, let us take a very short look at French. In French, no clitic doubling of any kind is attested, as shown in (20) with a [+human] DO.

(20) *French examples on clitic doubling*

- (a) L' enfant salue la professeur ('The child greets the_{FEM}&<SG> teacher')
- (b) L' enfant la salue ('The child her greets')
- (c) *L' enfant la salue la professeur
('The child her greets the_{FEM}&<SG> teacher')

In fact, no one has discussed the question why clitic doubling is out in French, since non-doubling would seem to be the normal and logical case.

A great difference between French and Spanish is the number of person-number forms of the tensed verb. As is well-known, French is not a null subject language, probably due to the poorness of the person-number affix. The verbal tensed paradigm is poor, if we compare it with the rich Spanish system. This might be the reason why the French clitics are not incorporated in the verb form as some kind of prefix: the suffix is on its way out.

Italian, however, has a reasonably great number of person-number affixes, but there is no clitic doubling.

In both languages, there is much difference in forms and the use of [-human] weak pronouns and of [+human], usually strong, pronouns.

6. Résumé

The intuitive basis of this sketch has been that the different treatment of human and non-human nouns, or rather, argumental phrases headed by nouns, in Spanish should be related to the gender of Spanish nouns, which is a formal property that is very much visible in the effects it has on matching phenomena. The gender of Spanish non-human nouns has no meaning, but it is very much part of the native speaker's linguistic knowledge. The gender of Spanish human nouns, on the other hand, is nearly always related with the sex of the referent, and it has meaning.

The hypothesis has been that the uninterpretable gender feature of non-human nouns blocks direct object clitic doubling and that the interpretable gender of human nouns makes direct object clitic doubling possible. The impossibility of using strong subject and object pronouns to refer to non-human referents can be linked to the uninterpretability of the gender feature of non-human Nouns.

An interesting final question is whether there are languages which have visible subject-verb or clitic-object agreement of an uninterpretable feature.

References

- Bosque, I. & Demonte, V. (eds). 1999. *Gramática descriptiva de la lengua española*. 3 Vols. Madrid: Espasa-Calpe.

- Cardinaletti, A. & Starke, M. 1999. The typology of structural deficiency: A case study of the three classes of pronouns. In *Clitics in the languages of Europe*, H. van Riemsdijk (ed.), 145–233. Berlin: Mouton de Gruyter.
- Chomsky, N. 1995. *The Minimalist Program*. Cambridge MA: The MIT Press.
- Chomsky, N. 2000. Minimalist inquiries. In *Step by Step. Essays on Minimalist Syntax in honor of Howard Lasnik*, R. Martin, D. Michaels & J. Uriagereka (eds). Cambridge MA: The MIT Press.
- Chomsky, N. 2001. Derivation by phase. In *Ken Hale. A life in language*, M. Kenstowicz (ed.). Cambridge MA: The MIT Press.
- Corver, N. & Delfitto, D. 1999. On the nature of pronoun movement. In *Clitics in the languages of Europe*, H. van Riemsdijk (ed.), 799–861. Berlin: Mouton de Gruyter.
- Fernández Soriano, O. (ed.). 1993. *Los pronombres átonos*. Madrid: Taurus.
- Fernández Soriano, O. 1999. El pronombre personal. Formas y distribuciones. Pronombres átonos y tónicos. In *Gramática descriptiva de la lengua española*, I. Bosque & V. Demonte (eds), Vol.1, Ch. 19, 1209–1273. Madrid: Espasa-Calpe.
- Gutiérrez-Rexach, J. 2001. Interface conditions and the semantics of argument clitics. In *Current issues in Spanish syntax and semantics*, J. Gutiérrez-Rexach & L. Silva-Villar (eds), 107–142. Berlin: Mouton de Gruyter.
- Jaeggli, O. 1993. Tres cuestiones en el estudio de los clíticos: El caso, los sintagmas nominales reduplicados y las extracciones. In *Los pronombres átonos*, O. Fernández Soriano (ed.), 141–173. Madrid: Espasa-Calpe.
- Mendikoetxea, A. 1993. Los clíticos como categoría subléxica de concordancia. In *Los pronombres átonos*, O. Fernández Soriano (ed.), Ch. 6, 205–230. Madrid: Espasa-Calpe.
- Ormazabal, J. & Romero, J. 2007. The object agreement constraint. *Natural Language and Linguistic Theory* 25: 315–347.
- Riemsdijk, H. van. (ed). 1999. *Clitics in the Languages of Europe*. Berlin: Mouton de Gruyter.
- Schroten, J. 2001a. The human factor in subject pronouns and possessive pronouns in Romance and Germanic languages. In *The expression of possession in Romance and Germanic languages*, L. Tasmowski (ed), 53–68. Cluj: Clusium.
- Schroten, J. 2001b. El género del nombre y su interpretación. *Moenia* 7: 113–122.
- Suñer, M. 1993. El papel de la concordancia en las construcciones de reduplicación de clíticos. In *Los pronombres átonos*, O. Fernández Soriano (ed), 174–204. Madrid: Taurus.
- Suñer, M. 1999. Clitic-doubling of strong pronouns in Spanish: An instance of object shift In *Grammatical analyses in Basque and Romance linguistics*, J. Franco, A. Landa & J. Martín (eds), 233–256. Amsterdam: John Benjamins.

Metalinguistic processing and acquisition within the MOGUL framework

Michael Sharwood Smith
Heriot-Watt University

MOGUL (Modular On-line Growth and Use of Language) is a framework that builds on Jackendoff's views about the language faculty. Jackendoff's approach not only has 'psychological' reality in that it claims to account for linguistic knowledge in the individual and the logical problem of language acquisition, it also has what might be called 'psycholinguistic' reality making claims about on-line language processing as well. This aspect is exploited in MOGUL's account of language acquisition and performance and, amongst other things, permits explicit accounts of metalinguistic ability, here claimed to be a mental system developed and operated largely outside the modular language system and resulting in a different type of grammatical knowledge, one which complements and also conflicts with intuitively cognised grammar to which we have no conscious access.

1. Introduction

In this chapter, I will present a view of grammar which essentially elaborates the claim that typical adult native speakers have access to not one but two parallel mental grammars of the same language. Both of the grammars are in some sense 'natural' but in other respects differ in a number of important ways. One of them is remarkably stable and relatively free of optional elements. This is the metalinguistic grammar which is represented in such a way as to allow conscious access and therefore reflection on grammatical properties and, in addition, the construction of a metalanguage to support this. The other grammar is grammar that is normally associated with Chomsky's long established term, 'competence' (Chomsky 1965). We can discuss both types of grammar in an abstract, time-free manner as well as in terms of real time processing, so psychological plausibility is not a way of distinguishing the two. Rather, it is the way in which they are represented in the cognitive system as a whole. In line with proposals by Sharwood Smith and Truscott (Truscott & Sharwood Smith 2004;

Sharwood Smith 2005¹), the chosen approach to be adapted for this purpose is Ray Jackendoff's (see also Carroll 2001 for a different extrapolation of Jackendoff architecture).

2. Psychological and psycholinguistic plausibility

When speaking of linguistic models or linguistic theories, 'psychologically real' is a descriptor that is usually meant to indicate an underlying claim about what such a model is attempting to explain. It does not suggest that the model is supposed to be a precise explanation of exactly how the mind works with respect to language. It does not have to describe the way language is processed on-line, for example. It does, however, suggest that the categories and mechanisms used to explicate properties of language at the very least facilitate the psychological explanation of linguistic performance or the way in which mental grammars are structured; this should be more than an adequate description of observed linguistic patterns; in some cases it may make a major contribution to the explanation of particular real time linguistic phenomena. For example, to say a phoneme is psychologically real is to say that some psycholinguistic aspects of speech behaviour are neatly captured by positing such a category. Chomskyan accounts of grammar can be called psychologically real in that they help us to understand the reason why small children are able to acquire complex linguistic systems at such a young age and to explain the knowledge underlying native speaker's grammatical intuitions, a knowledge that goes well beyond any explicit understanding they have about the structure of their mother tongue and for which the evidence in the environment can be lacking, elusive or ambiguous. Many linguistic theories in the past have had no such ambitions: there have been many grammarians over the centuries whose descriptions were not intended to provide insights into the human mind. At the same time, psychological reality seems to be a rather flexible concept. It is, at least, a relative one. One might claim that, say, lexical functional grammar or construction grammar are 'more psychologically real' than Chomsky's in that their formalisms aspire to be closer to the formalisms required to explain real time processes. In this respect then, it might do to distinguish between the more general, abstract claim of psychological reality, that is, that the categories and mechanisms used to explicate properties of language facilitate the psychological explanation of linguistic performance or the way in which mental grammars are structured and one

1. The theoretical framework to be used in MOGUL (Modular On-line Growth and Use of Language) elaborated in Truscott and Sharwood Smith (2004) and thereafter in a number of different publications focusing on different issues in acquisition theory, for example Sharwood Smith (2005); Sharwood Smith and Truscott (2005, 2006); and Truscott (2007). Not much will be said here about the developmental aspects of metalinguistic knowledge (see Sharwood Smith 2005 for a fuller account).

which, more ambitious still, claims 'psycholinguistic' reality where formalism in grammar are more obviously useful when coming to describe psycholinguistic mechanisms that operate in real time. By the same token, 'neurolinguistically real' would refer to grammatical formalisms that contribute directly to the description of neural mechanisms associated with language use. In these latter cases the term 'plausible' instead of 'real' is sometimes used. If something is not 'real' or 'plausible', it means that whatever else the model says, a particular aspect of it is unhelpful or counterintuitive in explaining given psychological phenomena. Thus connectionism may have, in some respects, a certain degree of neurolinguistic plausibility but a particular aspect of it might still be quite implausible, hence limiting, perhaps quite substantially, but not in itself negating its neurolinguistic plausibility in general. Chomsky and Miller, when they considered the possibility that syntactic transformations might be, in current terms, 'psycholinguistically plausible', were going beyond the more general claims about the psychological reality of generative grammar (Miller & Chomsky 1963; Phillips 2004). When it turned out that derivational complexity did not help to predict language processing facts, this had no impact on the psychological reality issue with respect to the aims and validity of generative grammar. It just meant that the syntactic transformations of that early version of generative grammar had little or no psycholinguistic plausibility. By way of contrast to this, the architecture of the language faculty as currently conceived by Ray Jackendoff is so designed that discussions of the abstract system can easily feed into discussion of real-time processing, as was already clear from his proposals, in the nineteen eighties, on language and computational mind, since, as he himself points out, they are clearly compatible with a number of psycholinguistic models used to guide experimental studies (Jackendoff 1987, 1997, 2002). The modules that comprise the architecture as described in strictly competence terms, for example the phonological module and the syntactic module, are matched by the modular systems that define how language is processed in real time. The phonological module in both types of description contain structural elements that are unique to phonology and cannot be integrated with syntactic elements either in some kind of derived structural representation with both syntactic and phonological properties nor in a configuration that relates to how this representation is constructed in real time. In both cases, the codes of the two systems are mutually incompatible. An interface system is required, in both levels of description, which links given phonological structures with given syntactic structures, but without any exchange of information between the modules, or any translation of one code into another. Put more the 'design' of the abstract linguistic system is clearly reflected in the structure of the parser. In one mode, we may speak of a phonological structure being 'indexed' with a syntactic structure, just indicating the relationship between them. In a processing, real time mode, we may talk of the selfsame structures, now activated in working memory, as being placed in registration with one another, or chained together, by the interface 'processor'.

The discussion that follows will focus on what is (by hypothesis) a different kind of linguistic knowledge exhibited by humans, namely metalinguistic knowledge, knowledge 'about' language as opposed to knowledge 'of' language and how this knowledge may be accounted for in a framework that purports to have a reasonable degree of psycholinguistic plausibility. Someone may have stable intuitions about the grammar of his or her language but may be quite unable to explain those intuitions, for example the fact that a particular utterance 'sounds' wrong' or 'sounds right'. As a matter of course, academic linguists will have much to say about the formal properties of utterances and they will have the metalanguage to formulate these views in a very detailed and explicit fashion. The average citizen will only have such terms as 'word' and 'sentence' and 'rhyme', and perhaps 'accent', 'syllable', 'vowel' and 'consonant' to help them talk about aspects of language. Metalinguistic ability also implicates consciousness, since by focusing our attention on the formal properties of language rather than the meanings conveyed, we become able to consciously manipulate them, as is the case when devising a play on words, a joke depending on a pun or thinking up words that rhyme or begin with a given letter, for example. As implied by the above, we all have some degree of metalinguistic ability ranging from a general awareness of language as an object to the very sophisticated levels shown by academic linguists, multilinguals, journalists, orators, novelists and poets. We owe much of our more sophisticated metalinguistic ability to literacy, that is, to the formal education that we have received, although metalinguistic awareness begins very early as we are acquiring our first language and is, it is claimed, particularly encouraged by the acquisition of more than one language (Ben-Zeev 1977; Bialystok and Majumdar 1998; Bialystok 2001; Hakuta & Bialystok 1994: 122). Although there has not been much in the literature on the relationship between an awareness of grammar and grammatical acquisition amongst very young learners, it has been a hot topic in second language acquisition research from its early days in the nineteen seventies. One of the reasons for this is that, with the increased metalinguistic knowledge gained during the course of school education, it seems natural to exploit it when teaching or learning a new language. Could knowledge about language and about the specific target language help to enhance or even substitute for the tacit knowledge gained subconsciously from primary linguistic data? The level of metalinguistic knowledge of the target language will vary with the individual learner and especially with the approach adopted by the teacher or teaching method. However, as observed by Krashen in the mid nineteen seventies, formal knowledge of the grammar appears to have a very limited effect on the growth of the second language (L2) grammar, the intuitive grammar as it were that underlies spontaneous, unreflecting performance, suggesting that there is no 'interface' between the two kinds of knowledge such that metalinguistic knowledge can affect subconscious knowledge of the L2 and that therefore telling learners about the formal properties of language was essentially a useless exercise. Krashen's particular views on this issue caused much heated discussion, but, although the field

of second language research from the eighties onwards turned its attention to applying generative grammar, mainly principles and parameters theory, to the task of explicating learner grammars, the general claim, rather than the specifics of Krashen's theory, that conscious learning has limited relevance has been more or less taken for granted by many researchers investigating the linguistic properties of learner grammars. This is true despite the long-standing debate about explicit and implicit knowledge, or declarative versus procedural knowledge that has been going on in more psychologically or psycholinguistically oriented circles (see for example Paradis 1993; Ullman 2001; and contributions by Robinson and De Keyser in Doughty & Long 2003). This discussion has not had much impact on the (especially generative) linguistic studies on second language acquisition. Since no framework has existed that provides the necessary the connections. This has created the bizarre situation of a rich literature on (first and) second language development with no apparent learning theory, i.e., a theory that does not just define the linguistic properties of various stages but explains the transitions between one stage and another (see discussion in Gregg 1993). At the same time, up until recently there has been no proper theoretical framework in which to formulate precise questions about the psychological and psycholinguistics status of metalinguistic knowledge. This is true, at least, in that there has not been much headway in specifying, in more than general terms, the role that metalinguistic knowledge has in first and second language acquisition; this one would like to have done in such a way that links up coherently with current views on linguistics and other relevant domains of cognitive science. Generative grammar in its most orthodox formulation has little to say about the status of metalinguistic knowledge except that it cannot be equated with competence or 'I-language'. Jackendoff's much wider ranging approach linking language-specific systems to cognition in general is attractive in this regard. For example, he is able to use his generative approach to language to develop explanations about musical ability (Lerdahl & Jackendoff 1983; Jackendoff & Lerdahl 2006). This concern about linking language with other aspects of cognition does, it turns out, make available a good platform on which to elaborate further the status of this 'non-intuitional' knowledge about language, a type of knowledge which is claimed to lie outside the language faculty. The discussion that follows briefly presents an elaboration of non-intuitional knowledge along these lines. This will be preceded by some general observations on noticing and awareness in language acquisition.

3. Noticing and awareness in L2

Turning now to the grammatical features of a second (or other) language (L2), the learner, in order to reproduce the correct patterns of the target language, must, logically speaking, first notice them. The precise nature of noticing is a big issue here but noticing, at least

in some sense of the term, must be a necessary condition for the patterns to appear later in the learner's own linguistic performance. Without noticing these language forms, it is logically impossible to make them part of the developing grammar of the L2. For example, you really do need to notice, in some sense, overt morphological forms that reflect syntactic agreement in L2 adjectives, that is, notice their phonetic and/or orthographic shape, and where and when they occur before you can hope to involve them in the way in which you, the learner, interpret L2 utterances and ultimately produce them yourself. This holds true whether or not learner performance is native-like. A learner with a native language that does not mark agreement on adjectives, for example, may use the wrong inflectional form on an L2 adjective in the wrong context but we can nevertheless be sure at least that the unfamiliar forms, i.e., agreement inflections, that appear in the speech and writing of native-speakers of that L2 have actually been, in some sense, 'noticed' by that L2 learner. Put more neutrally, somehow the forms have been 'registered' and 'retained' in some form or other and this has subsequently brought about changes in that learner's L2 performance (see discussion in Tomlin & Villa 1994). It is therefore beyond dispute that for any learning to take place, forms have to be perceived, processed in some way and retained, again in some form. In particular, where there is a mismatch between the learner's L1 and their L2, the existence of some novel, unfamiliar form or construction in the learners' L2 will provide us with convincing evidence of this because it cannot logically have been copied over from the L1. Put another way the appearance of what must be an unfamiliar form or construction, an inflectional morpheme or a determiner, say, that has no L1 equivalent in a learner's L2 repertoire presupposes some external input unless one espouses the totally implausible view that all aspects of all existing languages are available to us from birth precluding the need for exposure to input from outside. It seems logical and straightforward, then, to say that that noticing, in some sense of the word, is needed for developing new linguistic structure.

However, the question of whether the use of the term, 'noticing' implies a conscious awareness of the forms in question remains an open and important issue. Schmidt's Noticing Hypothesis defines the act of noticing as noticing surface elements but without any degree of analysis or understanding (Schmidt 1990, 2000). Most people accept that we, i.e. our minds, are able to register the presence of objects and events in our immediate environment without our necessarily being very aware of them. It is also uncontroversial to claim that awareness is a relative concept beginning with no awareness at all, or at most a very vague and transitory awareness that remains in short term memory for such a short time as not to trigger any reflection as to its nature or cause of the perceived object or event. The other end of the spectrum is represented by a heightened, intense awareness, the latter state requiring focal rather than peripheral attention. This type of awareness is associated with introspection, a 'conscious' mental activity that is characterised as 'metalinguistic' and which requires a heightened and sustained degree

of awareness to allow the learner to reflect on whatever they have noticed. For this to happen, the object or event must remain in memory for more than a few milliseconds and in considerably greater detail than is the case where only fleeting peripheral attention is implicated. In brief, then, the act of noticing a linguistic form or pattern may be associated with more or less awareness. Heightened awareness, however, does not necessarily involve a high degree of conscious analysis. Our ability to reflect on such forms and patterns in an analytic fashion requires a structured understanding of the properties of those linguistic patterns, something we can term 'metalinguistic knowledge' (Sharwood Smith 1994: 11). It also involves the possibility of registering gaps, that is, elements that are not on the surface but which are expected, what Chomsky has called indirect negative evidence (Chomsky 1981: 8–9). It seems plausible to assume that a learner is more likely to notice, say, the absence of a subject pronoun where in his or her L1 it would be expected, that is, when he or she is consciously reflecting on the properties of forms they are being exposed to. And many learners are less likely to notice this absence when they are processing the language input with little or no conscious awareness. This is of course an empirical question but the onus of proof must surely lie on those who wish to claim that gaps (i.e. expected but absent forms) are regularly registered subconsciously and are incorporated into the L2 grammar.

4. Metalinguistic and intuitional knowledge

The distinction between knowledge that we have conscious access to and therefore can reflect upon and knowledge that is intuitive, often called tacit knowledge, has been the subject of much debate in philosophical as well as linguistic circles. Stich (1978; see discussion in Schwartz 1986) refers to tacit knowledge as 'subdoxastic'. The two defining characteristics of subdoxastic knowledge are (1) its contents are not accessible to conscious thought and (2) they are 'inferentially insulated', that is, they have little inferential connections with the speaker's beliefs and so in principle can be in conflict with those beliefs. How could this be interpreted within a linguistic context? An example of such a conflict would perhaps be when one has a positive intuitive judgement about the grammaticality of a sentence that conflicts with a consciously held negative opinion about it as reflecting its 'bad grammar'. Tests of grammatical acceptability can therefore become questionable due to the unreliability of informant judgements. An informant might well intuitively know that a given sentence was grammatically acceptable in the sense that it was a possible sentence in that language and yet deliver a negative judgement because the construction is deemed to be unacceptable for reasons that are irrelevant to the researcher but which reflect linguistic prejudice, feelings about clumsy style, low probability or non-standard usage. Grammatical acceptability can mean different things to the researcher and the informant. Informants may find it impossible

or at least extremely difficult under certain circumstances not to insist on rejecting a perfectly 'grammatical' sentence even when the precise intentions of the researcher are explained to them. A more mundane example, often experienced by non-native speakers of a language, would be the tendency of native speakers to reject utterances as non-native which they, the native-speakers, without apparently being aware of it, regularly use themselves.

Whereas generative linguistics has elaborated on the tacit knowledge of language (or more properly of 'grammar' in quite a narrow sense) in great detail and in many ways, metalinguistic knowledge has been set aside and has rather attracted the attention of those more specifically interested in language development, and most especially in second language acquisition by adolescents and adults. For the early years of life, while they are busy becoming little native speakers, children only possess a rudimentary awareness of the formal properties of language although it has to be said that researchers in bilingualism have noted the relative superiority in this regard of bilingual as opposed to monolingual children (Hakuta & Bialystok 1994).

5. Metalinguistic learning: Its impact on language acquisition

Formal language teaching, in its most conservative form, the grammar-translation method, but also less radical versions where, alongside occasional attention paid to formal grammar, there is a general emphasis on completing communicative tasks, is characterised by its metalinguistic bias: learners are required to actively reflect on the formal properties of the language and may be asked to notice and memorise grammatical patterns, rules and exceptions. Explicit grammar teaching involves making learners aware of grammatical properties and explaining rules. With the knowledge that has been acquired as a result of such teaching, the learner is then much more able to notice and analyse examples of rules and violations of rules in their own or other people's utterances so that they ought to be able to correct grammatical errors. In addition to this, they should be able to give some explanation of why the errors are indeed errors. For this, they will need a metalanguage, in other words, a certain amount of terminology to make their explanations explicit. What is not guaranteed is whether they will be able to learn from the errors that they have noticed, analysed and occasionally explained to others, as a result. It remains a possibility that all this knowledge equips you with is the ability to explain and correct errors and not to stop making them in the first place. For people who think otherwise, it may come as a shock that, already at the close of the nineteen sixties, some researchers began to be convinced that nothing can be further from the truth. Forms that seemed to stare a learner in the face by virtue of their frequency and simplicity or because a teacher or textbook had deliberately tried to make them especially noticeable, would mysteriously refuse to appear in the

learner's own spontaneous performance: the early appearance of other native-like L2 forms not so frequent or not so salient in other respects, was, by the same token, also a mystery. Furthermore, after several decades of empirical research, no real consensus has arisen on this issue and the weight of the evidence continues to support the lack of effectiveness of grammatical instruction (Truscott 1998; see also discussion in Norris & Ortega 2000; and Doughty 2003). All we can agree on is that older learners who are understandably curious about their errors and want and expect explanations will be satisfied to receive some formal instruction about the grammar of the language they are learning and especially the correction, at least from time to time, of their grammatical errors and, furthermore, they will be frustrated if they do not get some explanation and correction. All of this does not take away from the fact that formal grammatical knowledge may still simply assist us to correct our grammar on the spot without preventing us from making precisely the same errors again and again in future. After all, the experience of learning in general from primary school onwards teaches us that we learn from making mistakes and being corrected. Why should language learning be any different? The possibility that there is in fact a difference is somewhat unpalatable.

Burt and Dulay's pioneering work on L2 developmental sequences inspired Krashen to compare immigrants who were only picking up English without any special assistance and those who were also undergoing formal instruction (see Dulay et al. 1982 for a general account of this research). He was interested in finding out whether the natural sequence in the acquisition of various morphemes in L1 acquirers of English was manifest in the performance of both groups of L2 acquirers and whether or not formal instruction brought any advantage. He found that on spontaneous tasks like picture description, learners from both groups produced the same kind of results already seen in the Burt and Dulay studies. The upshot of this experimental research suggested that L1 influence did not seem to affect the difficulty orders in the data and that formally instructed learners only produced different, and better results on the tasks that gave them time to reflect and self-correct and, it was claimed, showed the influence of their formally gained grammatical knowledge. This finding led to Krashen's well-known distinction between the process he called acquisition (subconscious learning) and the process he called (conscious) 'learning' (Krashen 1985). In Krashen's view, each process produced an entirely different cognitive outcome. Krashen's 'acquisition' was a natural process that caused grammars to grow beneath the level of conscious awareness while his 'learning' resulted in a more artificial, technical kind of learning (here called metalinguistic knowledge) which could only be used in limited ways. Furthermore, learned knowledge had no effect on the natural course of development. Subject P, for example, was consistently able to correct her own errors when confronted with them, produced virtually no errors in writing but consistently made errors in her spontaneous speech thus demonstrating, in Krashen's view, the ineffectiveness of relevant learned knowledge

to transform itself, over time, into acquired knowledge (Krashen & Pon 1977) although why she remained at this level and why her development did not eventually respond to the input via the subconscious channel was not explained.

The bottom line would appear to be that the role and psycholinguistic status of metalinguistic knowledge, linguistic knowledge that is 'learned' in Krashen's sense (see also Schwartz 1999), has long remained defined only in general terms. Krashen provided a rudimentary psycholinguistic basis for it while generative linguistically oriented discussions with few exceptions have largely focused on properties of tacit knowledge alone, formulating them with much greater sophistication than had hitherto been the case in second language acquisition studies but leaving metalinguistic aspects largely aside for others to sort out. In order to elaborate on the tacit/'meta' distinction with a comparable degree of explicitness, however, it is necessary to make use of a theoretical framework that allows both kinds of linguistic knowledge to be explicitly and coherently explicated, namely MOGUL.

6. Jackendoff's architecture: An overview

What follows immediately below is an attempt to provide the barest of sketches of, and then extrapolate from, Jackendoff's model, the first part of which must rely on the extensive argumentation and illustration in Jackendoff's own work (see 1987, 1997, and 2002).

Jackendoff's model requires us to distinguish between the core language system (not his term but indicating the modularised phonological and syntactic systems described above) and external systems that interact with it: there is nothing particular about this distinction in itself but the advantage of his model is that it enables much more specific things to be said about the interaction between language in the narrow sense, tacit competence, and those aspects of cognition implicated in the notion of language in the broad sense, i.e. anything studied in various branches of linguistics outside the core language system including much of semantics, pragmatics, style, register, language varieties, etc., etc., (c.f. Hauser et al. 2002). Moreover, it translates easily into a real-time processing system which is also the mode in which it will now be discussed here. The battery of autonomous modules that make up the core language system covers both phonology and syntax. This includes both processors that integrate structures from either the phonological and syntactic 'sublexicons', (in processing terms, memories) to build, respectively, phonological or syntactic structures on-line for both production and comprehension as also processors that mediate between the stores of structures, linking phonological structures with specific syntactic structures, and between those structures and the peripheral systems outside, for example the system that creates speech or that which creates meanings expressed in terms of conceptual structure. Conceptual

structure (broadly speaking, linguistically relevant meaning, i.e. semantic, pragmatic and discourse structure) lies outside the core system except the interface system that mediates between it and structures within the core system such as the SS-CS (syntactic structure-conceptual structure) interface that links specific syntactic structures with specific conceptual structures outside the core system. A traditional lexical item, for instance, does not exist as such, as a single unit, but is a chain of structures each produced by a different subsystem, *cat* being a PS (phonological structure), placed in registration with a SS (syntactic structure: N [sing] etc.) which in turn is placed in registration with a given CS (conceptual structure: informally, SMALL DOMESTIC ANIMAL etc.). These three structures can be seen as being indexed together in a manner familiar from standard generative grammar: $PS^i \leftrightarrow SS^i \leftrightarrow CS^i$. The unification or 'chaining' of these three quite distinct structures give us the equivalent of a lexical item in the more usual sense.

A crucial aspect of Jackendoff's modular approach is that phonological and syntactic modules are 'structure-constrained', hence the need for interface systems to associate one structure with another. Phonology doesn't 'speak' syntax and vice versa. You cannot have features from one system incorporated into structures from another system. And if systems within the core language system cannot talk to one another directly, nor can they talk with anything outside. Their codes are not translatable and are therefore mutually incompatible: no exchange of information is possible. The same goes for conceptual structure outside the core language system. This means that the quasi-lexical item exemplified by the $PS \leftrightarrow SS \leftrightarrow CS$ chain for 'cat' (' \leftrightarrow ' indicating the interface systems between the three modules) is just that, a chain of associated structures; it is not a single combined, integrated structure. You cannot merge phonological, syntactic and conceptual features into a single unit. Interface systems only link (co-indexed) structures; they do not convert one into terms of the other or into some shared generic system: the three codes remain distinct. In addition, what is different about conceptual structure, that is, apart from the fact that it lies outside the core language system, is that, even though the fine detail of its structural properties remain hidden from conscious inspection, its contents are available to conscious awareness. Conceptual structure is the basis for thought processes, the mind's 'lingua franca'.

In this way, we can conceive of tacit grammar, Krashen's subconscious 'acquired' knowledge, as sealed off from conscious processes using Jackendoff's model to make this insight much more explicit. Tacit grammar then becomes the informationally encapsulated phonological and syntactic systems, as described by Jackendoff, along with their specific lexicons (linguistic memories) and their processors, including associated interface systems. More particularly, we can conceive of metalinguistic knowledge as created *outside* tacit grammar as described above and built up via conscious processes based on conceptual structure. In short, metalinguistic grammar has no direct involvement with the structural system we call tacit grammar and it exists outside the sealed-off modular system in which tacit knowledge is created.

Some more things then remain to be done before we can see how the two grammars fit into the cognitive system as a whole. Firstly, it is necessary to define how this external knowledge about language and tacit knowledge co-exist in the same mind and especially how they are used in real-time processing. This can be done more or less within the terms of Jackendoff's ideas as developed originally in his 1987 book albeit in the manner proposed by Sharwood Smith and Truscott as part of their MOGUL framework (Truscott & Sharwood Smith 2004). Furthermore, it is necessary to say something coherent about how metalinguistic knowledge is formed, how metalinguistic learning differs from the acquisition of the tacit grammatical system. Here it is necessary to make reference to an important feature of MOGUL, namely Acquisition by Processing Theory (APT) which is a rather particular extrapolation of Jackendoff's (cf. Carroll 2001 for an alternative approach to development using Jackendoff's architecture) although space limitations mean that this developmental angle will not be elaborated on here in great detail (see Truscott & Sharwood Smith 2004; and Sharwood Smith 2007).

Essentially, and very briefly, APT removes the need for any domain-specific learning mechanisms (like a Language Acquisition Device) but, notably, without affecting domain-specificity itself. The act of processing means placing elements in some area of working memory. Memory is not a single general facility but domain-specific. Any such element is retained for anything from a millisecond (or fraction thereof) to a lifetime, hence almost instant attrition may often be the instant follow-up to acquisition in this minimal sense, or, depending upon circumstances, the appearance of elements in working memory may also be the precursor to a lifetime existence in memory. Events retained in memory have a resting level of activation which is a function of use (in processing). In this sense, the more something is involved in processing, the greater its chances of being retained for a longer period. The higher the resting level has become (or the lower its threshold in Paradis' terms; see, for example, Paradis 1993), the more chance it has of being chosen for participating, later, in the on-line construction of a representation, since processing of any kind, including grammatical parsing, is a matter of constant and intense competition of activated elements to be chosen for incorporation in an utterance or an interpretation of an utterance. Since we are dealing with modules with codes that are unique to each module, i.e. structurally constrained, and since each working memory is specific to that module, the way representations are built up or lost over time, i.e., during longer term acquisition or attrition, modular development is constrained by the principles of the module concerned so that there may be highly frequent elements in the learner's environment but these will not automatically be acquired because of their high frequency if the internal conditions are not right. So, within this framework, although consciously accessible grammatical knowledge, i.e. metagrammatical knowledge, and tacit, intuitional grammatical knowledge have a quite different epistemological and cognitive status and handled in quite different ways within the

cognitive system as a whole, they are nevertheless acquired and lost according to exactly the same principles, i.e. APT.

7. Metagrammar linking conceptual and auditory structures

Language, as every introductory course to linguistics makes clear, is a way of associating sounds with meaning. Perhaps we should replace the completely fuzzy term 'language' with a broad definition of 'grammar', let's say including phonology, morphology and syntax and some form of semantics. This holds equally for grammars that we build up via conscious reflection and manipulation. The metagrammar, be it fragmentary or sophisticated, must also systematically associate sounds with meanings, all of which must be accessible to conscious introspection but, in this case, what are the sounds and what are the meanings? Staying faithful to the principle of structure-constrained modules, the metagrammar cannot rely on phonological structures to generate sounds (words, phrases etc.) that we can be consciously aware of and can manipulate. The sounds it has recourse to are the auditory structures of words as opposed to their phonological structure. An easy way to appreciate the difference is to consider the effect of a completely unfamiliar foreign word. The parser will attempt, automatically, some kind of phonological interpretation using the most accessible resources at hand, typically L1 phonology, but we must assume that it is alien enough for the hearer not to be able to assign any meaning or morphosyntactic structure to it. Nevertheless, it is a sound and, if we are skilful enough at imitating it, we may well be able to attempt to reproduce it in the same way as we might mimic a non-linguistic sound.² Both linguistic and non linguistic sound patterns have auditory structure. In the case of linguistic sounds, there will always be an attempt to assign linguistic structure to them, initially phonological structure. That is because language processing is mandatory (in the way described in Fodor 1983). This means that as, say, a single word is acquired, a link is forged between the word's auditory structure and whatever phonological structure the parser is able

2. Jackendoff (1987) points out that we can report having had dreams in another language indicating that some properties of speech are clearly available to awareness. Also, in tip-of-the-tongue phenomena we are aware of certain properties such as how many syllables the word has, and what sound the word began with. Admittedly 'syllable' is a phonological (prosodic) category but syllables clearly have auditory counterparts which are easily recognisable alongside door knocks, drum beats etc. In other words, the acoustic properties of syllables will have some auditory representation (informally put, they will be audible as sounds) aside from their linguistic properties which are computed by the phonological processors.

to assign it.³ When we become aware of the sound of a word, when we ‘hear it in our head’, as it were, we are accessing not its phonology but its auditory structure. True, the word’s phonological structure will be activated at the same time, automatically, since there will be an interface between the two systems although phonology cannot respond to any information from the auditory structure as a result of our conscious manipulations of words which we can hear in our heads. Auditory code simply cannot be processed in the phonological module so conscious manipulation of phonology remains impossibility. So, when we indulge in such metalinguistic reflection as ‘the word ‘cat’ is a noun referring to a small domestic, feline animal”, we are consciously aware of the individual words thanks to their auditory structure and also to the concepts involved, thanks to their conceptual structure. In the case of more sophisticated linguistic knowledge, we will also have developed appropriate terminology like ‘noun’. In sum, then, we are able to conceptualise grammatical relationships and apply them to the linguistic units without any direct access to the grammar ‘within’, i.e. the tacit grammatical system. In this way, we might come to have beliefs about the tacit grammatical system that are in conflict with reality. We will also be able to put positive or negative values on particular constructions which we might term ‘good grammar’ or ‘bad grammar’, all of this as a result of conscious, reflective processes that take place outside the language modules, i.e. what has been called here the core language system. Using Jackendoff’s architecture, we can be more precise about the way the metagrammar of English can exist and be used in the same mind that has tacit grammar of the same language. It is also noteworthy that the metagrammar can be consciously moulded into a conveniently simple and rigid system to suit the requirements of prescriptive grammar whereas tacit grammar (phonology and syntax) is a law unto itself and contains more optionality, vagueness and in some respects the ability to undergo change during a users’ lifetime without the user being aware of it or even desiring it. The sheer inaccessibility of the core language system grammar permits this to happen. On the other hand, a very large part of language in the broader sense, having to do with all the semantic and pragmatic aspects of linguistic usage, falls outside the core language system and is couched in conceptual structure hence falling to an area of cognition that is accessible to conscious reflection and manipulation. In other words, the conscious use of language is what marks out the literate, educated user of the language and is the outcome of many years of education starting after the core language system has mostly

3. The more the phonology of a word is acquired certain phonetic properties (duration, for example) will be made more salient; these properties will necessarily be present in the auditory structure of the word but not necessary in a way that will allow the learner to be consciously aware of them. It is precisely phonological development that makes particular auditory properties salient.

been acquired. Metalinguistic knowledge can thus be situated alongside all other kinds of knowledge accumulated during one's school days and beyond.

8. Metagrammatical learning and metagrammatical fluency

Krashen's contention was that learned linguistic knowledge was very restricted: the average learner, unless s/he is an academic linguist, will only have recourse to a small number of linguistic rules and principles which can be used to consciously monitor and, where possible, patch up and improve the language utterances they are producing. Furthermore, the application of conscious grammar is virtually impossible in spontaneous speech, in Krashen's view, thus restricting its use to situations where there is time and the learner has opted to focus on the formal properties of the language and not just the message to be conveyed (Krashen 1985). The present view is somewhat different from Krashen's, particularly with regard to restrictions on use. As regards the tacit grammar, this must indeed proceed within constraints imposed by the principles that dictate the architecture of the language system in its narrow sense, i.e. universal grammar. What happens outside that domain is another matter. Metagrammar, like all metalinguistic knowledge, is free of these constraints and this has, following learnability theory, two consequences. First, the learner requires and can respond to correction and other kinds of negative evidence. If some construction is not instantiated in the language data and this has somehow been noticed (in the conscious sense of noticing) then the metalinguistic rule or principle involved can accordingly be amended to match the evidence.⁴ Secondly, there is no reason in principle why metalinguistic ability, the skill with which metaknowledge of language can be deployed during language performance, cannot reach high levels of fluency as is the case in the first language, say, with skilful public speakers. There is, then, good reason to doubt Krashen's radical, 'dysfluent' characterisation of learned knowledge which severely limits the ability of skilled second language learners to think on the fly and control their performance with what they consciously know about the second language.

9. Implications

Adopting a MOGUL perspective on metalinguistic knowledge has implications in a number of areas, including research methodology. Accepting the perspective on the

4. For further discussion of the distinction between grammatical and metagrammatical development, see Sharwood Smith (2002).

status and operation of metagrammatical knowledge that this framework provides, does not make the job of the researcher much easier as far as the interpretation of L2 performance data is concerned: rather the opposite. It is certainly not possible to avoid the patching up or “contamination” of the data by metalinguistic phenomena by just looking at spontaneous data alone and treating this as an adequate reflection of tacit grammatical knowledge, as Krashen has argued.⁵ Metalinguistic processing can be fast and efficient too. Metalinguistic knowledge can also be regarded as a form of literacy that is variably possessed by L1 and L2 speakers alike although L2 speakers may use it for additional purposes. Just as native speakers can be fluently literate, so too can L2 speakers be fluently metagrammatical. Not only does this give metagrammar a much more general and important role than one might initially ascribe to it: it is certainly not confined to L2 speakers; it also has a direct impact on our understanding of (apparent) L2 ultimate attainment,⁶ since native speakers may well be able to act like full native speakers and yet be deploying both kinds of grammar in their spontaneous performance. This does not exclude the possibility of native-like performance unassisted in this metalinguistic way but it does make the straightforward interpretation of linguistic patterns appearing in fluent performance even more difficult. Perhaps more to the point, in the context of the present discussion, a MOGUL perspective allows a much more fine-grained account of what metalinguistic knowledge is and how it fits into a general account of language cognition in the broadest possible sense.

References

- Ben-Zeev, S. 1977. The influence of bilingualism on cognitive strategy and cognitive development. *Child Development* 48: 1009–18.
- Bialystok, E. 2001. Metalinguistic aspects of bilingual processing. *Annual Review of Applied Linguistics* 21: 169–181.
- Bialystok, E. & Majumder, S. 1998. The relationship between bilingualism and the development of cognitive processes in problem solving. *Applied Psycholinguistics* 19: 69–85.
- Carroll, S.E. 2001. *Input and Evidence: The Raw Material of Second Language Acquisition* [Language Acquisition & Language Disorders 25] Amsterdam: John Benjamins.
- Chomsky, N. 1965. *Aspects of the Theory of Syntax*. Cambridge MA: The MIT Press.
- Chomsky, N. 1981. *Lectures on Government and Binding*. Dordrecht: Foris.

5. Felix has a different account which very broadly accords with the one presented here (Felix 1988).

6. Leaving aside the vexed question of what ultimate attainment actually means for multilinguals.

- DeKeyser, R. 2003. Implicit and explicit learning. In *The Handbook of Second Language Acquisition*, C. Doughty & M.H. Long (eds), 313–348. Oxford: Blackwell.
- Doughty, C. 2003. Instructed SLA, constraints, compensation and enhancement. In *The Handbook of Second Language Acquisition*, C. Doughty & M.H. Long (eds), 256–310. Oxford: Blackwell.
- Doughty, C. & Long, M.H. (eds). 2003. *The Handbook of Second Language Acquisition*. Oxford: Blackwell.
- Dulay, H.C., Burt, M. & Krashen, S. 1982. *Language Two*. Rowley MA: Newbury House.
- Felix, S. 1988. UG-generated knowledge in adult second language acquisition. In *Linguistic Theory in Second Language Acquisition*, S. Flynn & W. O'Neil (eds), 277–294. Dordrecht: Kluwer.
- Fodor, J.A. 1983. *The Modularity of Mind: An Essay on Faculty Psychology*. Cambridge MA: The MIT Press.
- Gregg, K. 1993. Taking explanation seriously; or, let a couple of flowers bloom. *Applied Linguistics* 14(3): 276–294.
- Hakuta, K. & Bialystok, E. 1994. *In Other Words: The Science and Psychology of Second-Language Acquisition*. New York NY: Basic Books.
- Hauser, M.D., Chomsky, N. & Fitch, W.T. 2002. The faculty of language: What is it, who has it, and how did it evolve? *Science* 298: 1569–1579.
- Jackendoff, R. 1987. *Consciousness and the Computational Mind*. Cambridge MA: The MIT Press.
- Jackendoff, R. 1997. *The Architecture of the Language Faculty*. Cambridge MA: The MIT Press.
- Jackendoff, R. 2002. *Foundations of Language*. Oxford: OUP.
- Jackendoff, R & Lerdahl, F. 2006. The capacity for music: What is it, and what's special about it? *Cognition* 100: 33–72.
- Krashen, S. 1985. *The Input Hypothesis: Issues and Implications*. New York NY: Longman.
- Krashen, S. & Pon, P. 1977. An error analysis of an advanced ESL learner: The importance of the monitor. *Working Papers in Bilingualism* 7: 125–129.
- Lerdahl, F. & Jackendoff, R. 1983. *A Generative Theory of Tonal Grammar*. Cambridge MA: The MIT Press.
- Miller, G. & Chomsky, N. 1963. Finitary models of language users. In *Handbook of Mathematical Psychology*, Vol. 2, R.D. Luce, R.R. Bush & E. Galanter (eds), 419–491. New York NY: John Wiley & Sons.
- Norris, J.M. & Ortega, L. 2000. Effectiveness of L2 instruction: A research synthesis and quantitative meta-analysis. *Language Learning* 50: 417–528.
- Paradis, M. 1993. Linguistic, psycholinguistic, and neurolinguistic aspects of 'interference' in bilingual speakers: The activation threshold hypothesis. *International Journal of Psycholinguistics* 9: 133–145.
- Phillips, C. 2004. Linguistics and linking problems. In *Developmental Language Disorders: From Phenotypes to Etiologies*, M. Rice & S. Warren (eds), 241–287. Mahwah NJ: Lawrence Erlbaum Associates.
- Robinson, P. 2003. Attention and memory during SLA. In *The Handbook of Second Language Acquisition*, C. Doughty & M.H. Long (eds), 631–678. Oxford: Blackwell.
- Schmidt, R. 1990. The role of consciousness in second language learning. *Applied Linguistics* 11(2): 17–46.
- Schmidt, R. 2000. Attention. In *Cognition and Second Language Instruction*, P. Robinson (ed.), 3–32. Cambridge: CUP.

- Schwartz, B.D. 1986. The epistemological status of second language acquisition. *Second Language Research* 2: 120–159.
- Schwartz, B.D. 1999. Let's make up your mind: 'Special nativist' perspectives on language, modularity of mind, and nonnative language acquisition. *Studies in Second Language Acquisition* 21: 635–655.
- Sharwood Smith, M. 1994. *Second Language Acquisition: Theoretical Foundations*. London: Longman.
- Sharwood Smith, M. 2002. Losing LAD and getting MAD: On Acquiring Parallel Language Systems. Paper presented at Department of Linguistics and English Language Colloquium series, Durham University, June 17, in England, UK. <<http://www.hw.ac.uk/langWWW/mogul/>>.
- Sharwood Smith, M. 2005. More on the meta mode: In search of deeper explanations for the role of consciousness in second language learning. In *Selected Papers on Theoretical and Applied Linguistics*, A. Psaltou-Joycey & M. Mattheoudakis (eds), 28–42. Thessaloniki: University Studio Press.
- Sharwood Smith, M. 2007. Revisiting the role of consciousness with MOGUL. In *Understanding Second Language Processes*, Z. Han (ed.), 1–15. Clevedon: Multilingual Matters.
- Sharwood Smith, M. & Truscott, J. 2005. Stages or continua in second language acquisition: A MOGUL solution. *Applied Linguistics* 22(2): 219–240.
- Sharwood Smith, M. & Truscott, J. 2006. Full transfer full access: A processing-oriented interpretation. In *Paths of Development in L1 and L2 acquisition*, S. Unsworth, T. Parodi, A. Sorace & M. Young-Scholten (eds), 201–206. Amsterdam: John Benjamins.
- Stich, S.P. 1978. Beliefs and subdoxastic states. *Philosophy of Science* 45(4): 499–518.
- Tomlin, R. & Villa, V. 1994. Attention in cognitive science and second language acquisition. *Studies in Second Language Acquisition* 16(2): 183.
- Truscott, J. 1998. Noticing in second language acquisition: A critical review. *Second Language Research* 14(2): 123–135.
- Truscott, J. 2007. Optionality in second language acquisition: A generative processing-orientated account. *International Review of Applied Linguistics* 44(4): 311–330.
- Truscott, J. & Sharwood Smith, M. 2004. Acquisition by processing: A modular perspective on language development. *Bilingualism: Language and Cognition* 7(1): 1–2.
- Ullman, M.T. 2001. The neural basis of lexicon and grammar in first and second language: The declarative/procedural model. *Bilingualism: Language and Cognition* 4(1): 105–122.

Catching heffalumps

Onset fricatives in child speech

Wim Zonneveld*
Utrecht University

This paper argues that a seemingly cognitive phenomenon, that of imitation in early child speech (the parroting of a model by the child), when approached in a formal manner, can provide unexpected insight into the nature of the developing grammar. With the help of a large corpus of – in principle – spontaneous speech of a single Dutch speaking child, it shows that it is highly rewarding to make a distinction between child utterances that are apparent direct imitations of a model, and unprompted utterances. A detailed discussion of empirical material illustrating the phenomenon of onset fricative avoidance leads to the conclusion that mere imitation is not just parroting, but is directly conditioned by the ambient grammar. The framework of Optimality Theory is used to demonstrate how, given a child grammar in which generally speaking markedness constraints precede faithfulness constraints, this finding can be formalised in terms of typically premature high-ranking faithfulness triggered by an imitation context. It is also shown that the formalisation of the analysis in terms of ranked constraints is compatible with independent O.T. theorising about the nature and development of early grammars.

1. Introduction

The aim of this paper is to contribute to the ongoing study of language acquisition, but with a twist: it makes an attempt at showing how a seemingly cognitive phenomenon, that of imitation (the parroting of a model by the child), when approached in a formal manner can provide unexpected insight into the nature of a developing grammar.

*I am grateful for comments and remarks first and foremost to Heleen Lohuis-Weber and Dominique Nouveau, who contributed significantly to its conception, and to Jan Don, Elan Dresher, Janet Grijzenhout, Neil Jacobs, Sandra Joppen, René Kager, Brigit van der Pas, Joe Pater, Daan Wissing, two anonymous referees, and the audiences at a variety of meetings where this paper was presented in different versions (including the universities of Utrecht, Düsseldorf, and Potchefstroom, South Africa). All responsibility for errors or misrepresentations is of course the author's.

In particular it aims to give an explanation of a specific kind of variation in child speech data for a single Dutch-speaking child, Nina, of which that between virtually perfect [1;11] *ofelant* for *olifant* ‘elephant’, and severely reduced but – strikingly – later [2;0] *fant* or *tant*, is an example. In the course of addressing such examples, and using a large corpus of what would normally simply be called spontaneous child speech, it will be shown that it is highly rewarding to make a distinction between child utterances that are apparent direct imitations of a model and unprovoked utterances. Doing so leads to the conclusion that mere imitation is not just parroting, but is directly conditioned by the existing child grammar. The framework of Optimality Theory will be used to demonstrate how this notion can be formalised in terms of (typically premature) high-ranking faithfulness. It will finally be shown how this formalisation is compatible with independent O.T. theorising about the nature and development of early grammars.

Following a considerable amount of literature of the past few decades, Lohuis-Weber and Zonneveld (1996) discussed the often lopsided behaviour of the sound class of fricative consonants in child speech. In the Dutch case they call attention to, that of Joost between the ages of [1;10] and [2;11], this child systematically – and in fact right until the very end of that period of investigation – failed to realise fricative consonants (in fact more generally continuants, so also glides and liquids) prevocally (in syllabic terms, in onsets), whereas he revealed no difficulties, from the very outset, producing them postvocally (in codas). Typical Joost data provided by these authors include [bus] for *bloes* ‘shirt, blouse’, [nes] for *fles* ‘bottle’, and [nis] for *televisie* ‘t.v. set’ at [1;10–2;2]; [naxnax] for *vrachtwagen* ‘lorry’ and [tejnɪs] for *televisie* at [2;2–6]; and [ojinat] for *olifant* ‘elephant’ [2;6]. (Notice that Joost unexpectedly but consistently replaces the fricative in question with the nasal non-continuant [n], a detail discussed but essentially left unexplained by Lohuis-Weber and Zonneveld, as it will be here). Observations about a similar positionally conditioned process can be found in a range of places in the literature, as in the references in the survey below.

1. Leopold’s (1947) diary study of Hildegard (bilingually acquiring English and German, investigated from [0;8] to [1;11]) includes the observation that she “appeared actively to avoid attempting fricatives”, and that her speech has an early stage at which initial fricatives are stopped, and [ʃ] is the only fricative in her phonetic repertoire, occurring finally representing all fricatives and affricates (even by metathesis, i.e. the transposition of consonants, as in [lɔɪs] for *story* (also see Ingram 1989: 376ff.)).
2. For Daniel at [1;4–2;1], Menn (1971) mentions output with fricative-free onsets such as [asʲ] for *watch*, [ɪsʲ] for *fish*, and [ay] for *slide*. Similarly, Waterson (1971) gives [ɪʃ] for *fish* and [ʊʃ] for *vest*.
3. Ferguson et al. (1973) observe that “fricative production emerges first in codas” (also see Ingram 1978, and Velleman 1996).

4. Gierut (1985/6) in a study of misarticulating children (investigated between the ages of [3;7] and [4;6]) notices as a possible pattern of misarticulation that “fricatives were only produced post-vocally, and never occurred word-initially; stops were produced instead”.
5. A report by Chiat (1989) on the behaviour of the misarticulating child Stephen [4;7–4;10] extensively discusses fricative stopping as specifically occurring in the onsets of syllables.
6. Leonard and McGregor (1991) point out examples of metathesis such as [uz] for *zoo*, [ops] for *soap*, and [ainf] for *fine* (also see Velleman 1996, and Bernhardt and Stemberger 1998: 432ff.).
7. In Jarmo’s (and other Dutch children’s) data, Fikkert (1994) detects three strategies of dealing with initial fricatives, before age 2: they are either deleted (e.g.: [ɪs] for *vis* ‘fish’), replaced with *h* (e.g.: [hisə] for *fietsə* ‘ride a bicycle’), or replaced with a plosive (e.g.: [plax] for *slak* ‘snail’).
8. Dinnsen (1996: 137) mentions Farwell (1976), Ferguson (1978), Stoel-Gammon and Cooper (1984), and Stoel-Gammon (1985) as “claim[ing] that, for at least some children, word-final position is a favored context for the acquisition of fricatives”.
9. Edwards (1996), in a study of the child Jason “with a developmental phonological disorder”, reports on tests conducted at [4;8] showing that he produced absolutely no fricatives in word-initial position, whereas “60% of word-final fricatives were correct”.
10. Grijzenhout and Joppen (1998) show that between [1;2] and [1;7] the German child Naomi has no fricatives in onsets whatsoever (cf. [Auba] for *sauber* ‘clean’, and [tu:a] for *Schuhe* ‘shoe’ at [1;4/6]), whereas she realises them word-finally (cf. [auf] for *auf* ‘on’ at [1;5]).
11. Bernhardt and Stemberger (1998: 432) observe, as a “common pattern” in child speech, that “[f]ricatives appear first in codas”, citing the reference in [6] above, and providing data from English child Jeremy, who uses “default insertion of coda [s]” (recall [1] above) as a response to an apparent prohibition against two non-continuant in a single output, as in [ɪʌs] for *milk* and [mɛʰs] for *bandaid* (in which, for independent reasons, the non-continuant nasal surfaces in the onset).

A number of properties are shared by these references and examples. First, they confirm that often in child speech failure to produce continuants in onsets is accompanied by the ability to produce them in codas. Second, the avoidance strategy is not language-specific, as the examples come from English, Dutch, and German. Third, there appear to be a number of ways of avoiding fricatives in onsets: to substitute them with a different consonant (as in Dutch [hisə] for *fietsə* and [plax] for *slak*); to omit them altogether (substituting them, as it were, with ‘zero’, as in English [ɪsʰ] for *fish* or Dutch [ɪs] for *vis*); or to move them to a preferred constituent (as strikingly in English [uz] for *zoo*).

Given these empirical preliminaries, the structure of this paper is as follows. Sections 2 and 3 provide case studies of onset fricative avoidance in the speech of two Dutch children, Joost (introduced above) between [1;8] and [2;11], and Nina between [1;3] and [2;3]. Joost's speech exhibits a relatively consistent substitution pattern, whereas Nina's exhibits much more variation, both synchronically and developmentally, and therefore poses the greater challenge. Section 4 has two goals: it shows how the accuracy of Nina's output (or, in terms of this paper's discussion below: its 'faithfulness') increases with time; and it introduces as a property of the data the distinction between a child's direct imitations immediately following an adult's model in a recorded exchange (such as – typically – looking at pictures in a picture book) on the one hand, and spontaneous speech on the other hand, whose utterances typically lack such a direct model. In the second half of this paper, this distinction leads to the association of unexpected cases of prompted imitations with premature faithfulness (Sections 4 and 5), and to the notion – only informally expressed in the child speech literature – that direct imitations are not just parroting, but can be conditioned by the existing child grammar (Sections 6 and 7). The actual analysis takes place within the framework of Optimality Theory, including the suggestion (Section 8) that its formalisation is compatible with O.T. ideas about the nature (and development) of early grammars.

2. Joost's treatment of fricatives

The data in (1) below, from Lohuis-Weber and Zonneveld (1996), illustrate the behaviour of continuant consonants, prominently including the fricatives, in Joost's speech between [1;10] and [2;11]. In the earliest stage the output is monosyllabic, and if the input is polysyllabic, the (rightmost) stressed syllable is selected towards meeting this condition. Later, output becomes progressively more faithful to the input. Under-scored *n*'s indicate the cases in which a target continuant is substituted by this nasal consonant.

(1) Continuants in Joost's speech [1;10–2;11])¹

a. [1;10–2;2] (rightmost) stressed syllable selected from the adult input

1;9	[tɛɪn]	trein	'train'	1;11	[<u>n</u> ɑm]	warm	'warm'
1;10	[<u>n</u> un]	schoen	'shoe'	1;11	[bus]	bloes	'shirt, blouse'

1. The Joost data of this section are based on the transcriptions of 65 recordings of approx. 15 minutes each, made between April 1990 and July 1991, assembled on 11 audiocassettes. Typical situations involved the child playing, reading, dressing, or helping out with groceries; the child's favourite reading material was Richard Scarry books, which helped to elicit many data.

1;10	[<u>n</u> es]	fles	'bottle'	1;11	[<u>n</u> ik]	muziek	'music'
1;10	[ox]	ógen	'eyes'	1;11	[<u>n</u> is]	tèlevísie	't.v. set'
1;10	[<u>n</u> ix]	líggen	'lie down'	1;11	[bof]	bóven	'upstairs'
1;10	[tɛix]	tijger	'tiger'	2;0	[<u>n</u> aj]	zwaar	'heavy'
1;10	[<u>n</u> œyj]	lúier	'nappie'	2;1	[<u>n</u> as]	wásen	'wash (oneself)'
1;10	[<u>n</u> ot]	ólifant	'elephant'	2;2	[p <u>a</u> n]	pánda	'panda'

b. [2;2–2;6]: stressed syllable(s) selected

2;2	[<u>n</u> aj <u>n</u> is]	wálvis	'whale'
2;3	[tej <u>n</u> is]	tèlevísie	't.v. set'
2;4	[oj <u>n</u> ot]	ólifant	'elephant'
2;4	[pek <u>n</u> as]	spèculáas	'ginger bread'
2;5	[<u>n</u> in <u>a</u> f]	giráf	'giraffe'

[2;6–2;11]: syllable-faithful output

2;7	[oj <u>n</u> ot]	ólifant	'elephant'
2;11	[tej <u>a</u> nisi]	tèlevísie	't.v. set'

Clearly continuants are replaced with *n* when in onsets, and preserved when in codas: the single item *fles* [nes], for instance, embodies both of this generalisation's branches. Some of the items, however, also raise the suspicion that occasionally independent processes interfere, affecting the generalisation's transparency. It is less than perfectly clear, for instance, why the output representing the target *televisie* (*tè-le-ví-sie* in adult syllabification) always shows just a single *n* no matter what its number of syllables, and the same holds for that representing *olifant* (*ó-li-fànt*). In Lohuis-Weber and Zonneveld's view, the major interfering factor in these data is the Stress-to-Weight Principle (SWP), the idea proposed in the literature (see below) that there is a strong tendency for stressed syllables to be typically heavy or become heavier, i.e. to 'acquire weight', for instance by attracting a consonant from the adjoining onset (also known as 'coda capture'). Assuming Joost's syllabifications to be *tèl-e-vís-ie* and *ól-i-fànt*, with non-adult stressed heavy syllables, selecting the main stressed syllable of the item *televisie*, for instance (which is *-ví-* in adult speech), results in child output [nis]; and selecting more syllables results in an output ending in [-nísi]. Moreover, the intervocalic *l*s of the target sequences *tèl-(vísie)* and *ól-i-(fànt)* are outputted as [-j-] rather than [-n-] because coda liquids generally appear as [-j], cf. [naj] for *zwaar* 'heavy' at [2;0]; [naj] for *schaar* 'scissors', [tɛj] for *ster* 'star' at [2;1], [nɔj] for *snor* 'moustache', [maj] for *allemaal* 'all of them', [dɪj] for *krokodil* 'crocodile' at [2;2]; and [hejikɔpt(əj)] for *hèlicópter* at [2;6–9]. Confirming evidence for the analysis is suggested by output such as (perfectly correct!) [kɔfi] for *kóffie* 'coffee' [2;7], and [pinazi] for *spinázie* 'spinach' [2;5], discussed in Lohuis-Weber and Zonneveld (1996) and Zonneveld (1999), in which fricatives survive when they immediately follow a stressed vowel.

SWP cases have been discussed for (adult) English in work such as Selkirk (1982) and Myers (1987), often under the heading of Resyllabification. Myers (1987: 496) visualises this process as below:

- (2) Resyllabification as in Myers (1987):
- | | | |
|---|--|---|
| x | | |
| x | | x |
| σ | | σ |
| \ | | ≠ |
| C | | |

These authors argue that a range of phonological processes of English are triggered (or blocked) in a configuration in which an apparent onset consonant appears to be captured in a coda. Among the processes discussed are (American English) flapping, aspiration, and *h*-deletion. Further discussion of the SWP can be found in Hammond (1992, 1999), Pater (1994), and Hammond and Dupoux (1996). Stemberger (1996: 73–4) describes another case from child speech: “It is often stated that intervocalic consonants should be onsets rather than codas: *bucket* [b.ʌk.ət], not [b.ʌk.ət]. [...] If an intervocalic consonant is in an onset, this seems to predict that the constraints on the features that may be present should be the same as in onsets and different from the constraints on codas. This is often incorrect. For all three of my children, for example, the intervocalic /k/ of *bucket* was realized as [k] when velars were allowed in codas but not in onsets. This implies that the intervocalic consonant was in a coda (where [C-dorsal] always survived in the output) rather than in an onset (where [C-dorsal] never survived).” The child speech variant of the process was discussed extensively in Bernhardt (ed., 2002), followed by an exchange between Marshall and Chiat (2003) (con coda capture) and Zonneveld et al. (2007) (pro) on the data originally presented in the reference in [5] above. Another variant of the acquisition of weight is by vowel lengthening, as discussed in Goldsmith (1990: 157ff.) for Chimalapa Zoque (Mexico), Crosswhite (1998) for Saipanese Chamorro (Micronesia), and Pearce (2006) for Kera (Eastern Chadic). Arias and Lleó (2007) observe on the acquisition of Spanish between [1;3] and [2;3] that these “children produce trochees target-like from very early on [by] lengthen[ing] the stressed syllable by resorting to a glide or an approximant. [This] we analyze as an effect of Stress-to-Weight, i.e. ‘if stressed, then heavy’”.² Joost’s data can

2. Prince (1990) argues in favour of SWP’s “converse”, i.e. the W(eight-to-)S(tress) P(rinciple), as universally applicable. He “specifically den[ies]” (p.358) the SWP the same status, reanalysing Myers’ 1987 central case of “resyllabification” (*sane* [seyɪn] vs. *sán-i-ty* [ˈsænɪti]) underlying English Trisyllabic Shortening: this is not vowel shortening in an excessively heavy syllable, but the result of a universal preference for the equal distribution of weight among the syllables of a trochaic foot. In a mollifying footnote, however, he allows for the possibility “that

be analysed as exemplifying such a case, too, using coda capture to add weight to the stressed syllable.

An intriguing final case is Joost's pronunciation of the town's name of *Utrecht*, whose output is [ytɛxt] (at [2;9]³), whereas clearly – based on the adult syllabification *U-trecht* - [ytɛxt] is the expected form, with onset cluster reduction as in [kɔk] for *klok* 'clock' at [1;9–2;3], [tɛɪn] for *trein* 'train' at [1;9–2;11], [bɔt] for *brood* 'bread' at [2;1–8], and [bʊmkɔj] for *bloemkool* 'cauliflower' at [2;2]. Highly interestingly, a completely unexpected output is derived here by a feeding role for SWP (cluster-final /-r/ becomes onset-initial and turns into [n-]) rather than the usual bleeding role (as typically in coda capture cases).

One of the most striking aspects of the Joost data is the extreme consistency exhibited in processes such as fricative (continuant) avoidance, SWP (coda capture), liquid avoidance in codas, and cluster reduction, throughout the period of investigation, i.e. between [1;8] and [2;11]. The clearest developmental pattern is that of the increase of number of syllables towards virtual faithfulness at [2;11], but the processes just mentioned show no sign of fading at that same point. This is unusual, and in fact Joost presents a case of delayed development (misarticulation at a late age, see [4] and [5] above), which had to be remedied by speech therapy. This special case of idealised data may be welcome to the analytical linguist in severely reducing the (potentially overwhelming) variation otherwise characteristic of child speech data, but also implies an element of idiosyncrasy, and eliminates the opportunity to use variation as an independent source of interest. This is the vein in which we turn to the more realistic Nina case which – as pointed out in the introduction – will supply the empirical underpinning of the central point of this contribution.

3. Nina's treatment of fricatives

Joost substitutes *n-* for continuants (fricatives and approximants) in onsets, and the previous section developed an analysis of the major components of that phenomenon. This section considers the speech of Nina, between the ages of [1;3] and [2;3], focusing on her treatment of fricatives (the approximants, i.e. liquids and glides, will be left out of consideration from this point onwards, as less relevant to the discussion's issues.). The exposition uses the transcriptions of approximately 60 hours of

Stress-to-Weight is a principle, but one with a different position than Weight-to-Stress in the ranking of rhythmic priorities”.

3. Joost was raised in the Bilthoven suburb of Utrecht, the university town at which his mother was a student.

taperecorded speech, from between the ages of [0;3] (July 1990) to [3;9] (January 1994), assembled on 46 audiocassettes. As an introduction, consider the three transcribed passages in (3).

- (3) Three passages from the Nina files
- a. [1;8] On not being allowed to handle a jar top at the dinner table.
 Nee niet doen Nina, da's vies. tis. [tis]
 Laat maar staan, laat maar staan. taan. [tan]
 [Please don't, Nina, that's dirty. Dirty. Leave it stand alone. Stand.]
 - b. [1;11] Looking at pictures in a picturebook
 Wat doen ze daar? temme.
 Zwemmen he? zemme.
 [What are they doing there? Swimming. Swimming, right? Swimming.]
 - c. [2;0] Looking at pictures in a picture book.
 Dit, wat is dit? vlag. [vɫax]
 Een vlag, goed zo. nee. [ne]
 En dit is een? n fant. [n fant]
 Een olifant. n tote flenk. [n totə flɛŋk]
 Een grote slurf, ja. uff. [œf:]
 [This, what is this? Flag. A flag, well done. No. And this is a...? An elephant.
 An elephant. A big trunk. A big trunk, yes. Trunk.]

A clear similarity between Joost's and Nina's data resides in the fricatives' lopsided distribution: in (3a) and (3b) fricatives are avoided in onsets ([tis, tan, tɛmə]), whereas throughout (3) they occur word-internally and –finally. A clear difference with Joost concerns the development from (3a) to (3c): (3a) contains (mere) attempted imitations of adult models adhering to the fricative constraint; on the other hand, (3b) and (3c) contain both successful (virtually faithful) and unsuccessful (less faithful) output. Example (3b) has [t]emme followed by [z]emme, and (3c) exhibits (no fewer than) three instances of fricative-initial output, to wit, [v]lag for vlag 'flag', truncated [f]ant for olifant 'elephant', and [f]lenk for slurf 'trunk' – accompanied by two cases of persistent onset-fricative avoidance: [t]ote for grote 'big', and [œf:] for slurf 'trunk'. In addition to this, a second kind of development takes place: in (3b) adult interference causes a crucially improved output form; therefore this passage also illustrates the child's ability to adjust output when prompted by a correct model. Passage (3c) shows a different, more sophisticated ability, namely to produce successful output – from the point of view of the fricative constraint – in the absence of a direct model. Thus, the successful output in (3b) seems qualitatively different from that in (3c); together with the other remarks of this paragraph, this suggestion will serve as the backdrop for the discussion of the Nina data in the remainder of this paper.

As established so far, like Joost Nina avoids output onset-fricatives at least in the initial stages of the period under investigation. She differs from Joost, however, in

employing a number of distinct avoidance strategies, the range of which is illustrated in (4). (Substitution by *n-*, it can be noted, is conspicuously absent.)

(4) Nina's avoidance strategies [1;3–2;0]

substitution in onset	1;8–10	[to]	zo	'like that'
	2;0	[te]	zee	'sea'
survival in coda:	1;8	[ox]	oog	'eye'
		[ʌuf(ə)]	elf [-l(ə)f]	'eleven'
		[mes]	mes	'knife'
stopping:	1;9	[as, ɛis]	ijs	'ice(-cream)'
	1;8	[bidada]	vanille [-nijə]	'vanilla'
	1;9	[bo]	viool	'violin'
stopping to dental:	1;10	[piŋə]	beginnen [-ə] ⁴	'to start'
	1;3–9	[tɪs]	vis	'fish'
		[tɪtə, tɪdə, tɪda]	zitten	'to sit'
	1;5	[tɔk]	vork	'fork'
	1;8	[tɛs, tɔsj]	zes	'six'
		[tɔs:]	gras	'grass'
		[dəsa]	flesje	'little bottle'
		[datə]	centen	'money,(cent) coins'
		[tataxs]	zaterdag	'Saturday'
	1;9	[tas, tɛs:]	fles	'bottle'
		[tœyt]	fruit	'fruit'
		[dux]	vroeg	'early'
		[taupo, taupf]	shampoo	'shampoo'
	1;9–10	[tits, dits]	fiets	'bicycle'
	1;9–11	[dɔk, tœ-ka]	varken	'pig'
stopping & harmony:	1;9	[koka]	sokken	'socks'
		[kɔkə]	zoeken	'to look for'
		[bopə]	schoppen	'to kick'
substitution by <i>h-</i> :	2;0	[peba]	zebra	'zebra'
	1;8	[hɛɪf, hauf]	vijf	'five'
		[hup]	soep	'soup'
		[hɪft]	schrift	'notebook'
deletion:	1;9	[hɔp]	schop	'spade'
	2;0	[hɪtə]	zitten	'to sit'
	1;8	[up]	soep	'soup'
		[əjə]	schande	'shame'
		[la]	vla	'custard'
	1;9	[est]	geest	'ghost'

4. In Dutch, the written ending *-en* is usually pronounced as *n*-less [-ə]; this will not be separately marked below.

As indicated by the number of examples per process, stopping to (typically a dental) plosive is by far the most frequent process. The passage in (3c) above, with its 3 fricative-initial items, suggests that at the age of [2;0] production of fricative-initial output is really starting to kick in. This is confirmed by the fact that most of the examples in (4) stem from [1;8–10] rather than later. In a separate and detailed investigation of Nina's fricative output, Van der Pas (1997: 21/43) gives the following numerological summary of these consonants' realisations in the course of the period of investigation:

(5)	Age	no. of targets	fricatives realised	fricatives avoided
	1;8	34	5 (15%)	28 (85%)
	1;9	67	16 (24%)	41 (76%)
	1;10	36	12 (33%)	24 (67%)
	1;11–2;0	53	29 (55%)	24 (45%)
	2;1	63	55 (87%)	8 (13%)
	2;2	140	135 (96%)	5 (4%)
	2;3	103	91 (83%)	12 (17%)

Nina's avoiding initial fricatives steadily declines with age, [2;0] being a hinge point, with a strikingly rapid growth towards accuracy at [2;1].⁵

Word-internally, Joost had coda-captured continuants surviving by the SWP, recall the contrast between *giráf* [ninaf] versus *kóffie* [kɔfi]. The picture emerging from the Nina data for non-initial fricatives is that of (6).

(6)	a.	Nina data: word-internal fricatives, post-stress		
	1;8	[neyə, na-kœn]	négen	'nine'
		[dœyva]	stúiver	'five cents coin'
		[taza]	mayonáise	'mayonnaise'
		[tafətəl ə]	verháal vertellen	'to tell a story'
	1;8–9	[kasa]	kássa	'cash desk'
		[tafəl, pafɔ]	táfel	'table'
	1;9	[dœf əl ə]	dolfijn	'dolphin'
		[dixœyx]	vliegtuig	'aeroplane'
		[bufar, tofnar]	tóvenaar	'magician'
		[aʊzəz]	(dino-)sáurus	'dinosaur'
	1;9–10	[tustərs, ɪsa]	lúcifers	'matches'
	1;9–11	[boʏɔ, toʏɔ, koʏɔ]	vógel	'bird'
	1;10	[ovə]	óver	'over'
		[hœysə]	húizen	'houses'
		[ɛsəmo]	éskimo	'eskimo'
		[tɪfɔ]	pantóffels	'slippers'

5. The data of [1;8] are not in Van der Pas, but have been retrieved from the same files. The slight 'regression' observed at age [2;3] may or may not be significant, but it will be left undiscussed here.

b. Nina data: word-internal fricatives, pre-stress

1;5–10	[dajt, pəpχaj, p-paj, taj]	papegáai	‘parrot’
1;8	[tywtet, mefekt]	perféct	‘perfect’
	[bɛlt]	Zónneveld	(family name)
1;8–10	[bebiton, petədon]	babyfóon	‘baby alarm’
1;9	[didətətədə sant]	croissánt	‘croissant’
	[ɣyzyzisi]	televisie	‘t.v. set’ (cf. fn.6)
1;10	[tafe, kɔfe]	café	‘cafe’
	[taɣent, tastent]	agént	‘policeman’
	[owəvar]	óoievaar	‘stork’
	[puxti]	spaghétti	‘spaghetti’

The data in (6a) clearly evoke Joost’s: fricatives do not appear in onsets (from the adult point of view),⁶ but they are used freely and frequently intervocally immediately after stressed vowels. Perfect output such as Joost’s [kɔfi] ‘coffee’ is matched by an equally perfect Nina item such as [kasa] ‘cash desk’ (which occurs no fewer than 4 times in this shape in the files). Before conclusions are prematurely jumped to, however, the (6b) data must be carefully scrutinised, too: what happens (word-internally) before stressed vowels? Here, the picture is much less clear: at [1;5] both [dajt] and [pəpχaj] are found, for instance, of which the former is a fully expected truncated form with dental plosive substitution, whereas the latter seems, at this early age, strangely complex and faithful. The same holds for a pair such as [tywtet] versus [mefekt] at [1;8]. At this point in the discussion these data are too few and inconsistent to allow a choice between two possibilities: either Nina follows Joost in resyllabifying fricatives (leaving some cases in (6b) as an unexplained residue), or she follows – perhaps hesitantly – a different pattern, one reported in the literature for English (for misarticulating Jason at age [4;8] in Edwards (1996), recall [9] in Section 1): no fricatives in *word*-initial onsets, and fricatives realised elsewhere regardless of the difference between post- and pre-stress. In Jason’s speech a frequent substitute for initial fricatives was [kj-], as in [kjʌm] *thumb*, [kjop] *soap* and [kju] *shoe* versus [maʊs] *mouth*, [tis] *teeth*, and [pʃwæf] *giraffe*. Word-internal fricatives survive, both before unstressed and (secondarily) stressed vowels, cf. [kʲavou] *shovel*, [kʲizouz] *scissors*, and [ɛufwɪnt] *elephant* next to [təʃom]

6. Curiously, at a very early stage fricatives were frequently produced in Nina’s speech in a highly specific sense: she was able to – and apparently delighted in – producing long stretches of speech containing chunks each starting with a velar fricative. The transcriptions for [1;3] contain sequences such as “gegegewew gegegewe geke ta m! Gegege geje dm gegewem geke ta m! Gegegegejegeje ge t ge ge m m ge ke!”, etc. Furthermore, at this same age adult attempts at eliciting responses by offering models, often met with a repertoire of similar ‘stock’ templatic answers, of a “gegugat”, “gegados”, “gegigaat”, or similar shape. This output, interesting though it may certainly be, was left out of consideration in the current overview.

for exchange (3b) provides the opportunity of confirming that – at this stage, for this item – the child in her speech (still) avoids initial fricatives, the correct consonant only appears when and because a direct model has been offered. (The same holds for the process of initial cluster reduction, as in *d(r)aaien* in (7b).) This methodological point is underscored by a juxtaposition of passages such as those below, in which the item *olifant* ‘elephant’ appears twice in the child’s speech, in different phonological shapes ((8b) repeats (3c)):

- (8) Partial transcriptions of Nina’s use of the item *olifant* ‘elephant’
- a. [1;11] direct ‘imitation’, looking at pictures in a picture book.
- ’t Boek is open, maar weet je wat voor dier dit is? ede ope.
Is dit een pinguin? neeee...
Oh, da’s geen pinguin, wat is ’t dan, hoe heet dit? pingin.
Olifant. ofəlant.
[The book is open, but do you know what kind of animal this is?
Is open. Is this a penguin? No-oo... So that’s not a penguin, then
what is it, what’s it called? Penguin. Elephant. Elephant.]
- b. [2;0] ‘spontaneous’ utterance, looking at pictures in a picture book.
- Dit, wat is dit? vlag.
Een vlag, goed zo. nee.
En dit is een? n fant.
Een olifant.
[This, what is this? Flag. A flag, well done. No. And this is a...?
An elephant. An elephant.]

In these two cases, which were mentioned earlier in the Introduction, the syllabically more faithful variant of *olifant* appears approximately one month before the truncated one. Such a seeming regression in faithfulness is of course interesting from a developmental point of view, but in assessing this case we should not fail to notice that the former form occurs in a direct imitation context, whereas the latter is uttered spontaneously: if anything, imitation forms can be naturally expected to be closer to the target than spontaneous forms. (Compared with truncated [f]ant, the imitation item [ofəlant] not only has more syllables, but interestingly exhibits metathesis to boot, i.e. the consonants /l/ and /f/ are transposed; this observation will be further addressed below.)

Surveying the full Nina fricative data available for stages before [2;0], adds credence to the distinction between imitations and spontaneous speech. As early as the age of [1;3], in the same session in which *vis* ‘fish’ is produced as [tɪs] and *zitten* ‘to sit’ as [tɪtə] (recall (4)), *fles* ‘bottle’ is produced as [foets] with fricatives occupying both syllable margins. The output [tɪtə] is an unprovoked reaction to a picture of a rocking horse, in the first interpretable spontaneously produced item of the data. Both other items directly follow models, implying that in the case of [foets] imitation has succeeded in forcing a high degree of faithfulness. At [1;5], [dajt], in a (7)-type exchange,

is a spontaneous reaction to a picture of a *papegaai* 'parrot', followed by – when corrected by the adult interviewer – much improved [papxaj].

It is only from [1;8] onwards that initial fricatives start appearing more frequently. At [1;8], in a collection of several dozens of relevant items (cf. (5)), only four cases incorporate an output initial fricative: [zɛʏəs] for *zeg* is 'say, please', [F]oka (a neighbour's name), [vɔxk] for *vork* 'fork', and a sequence in which the emphatically presented model adjective *perfect* first leads to [tywtɛt] and only in the second instance to [me-fɛkt]; in fact, in each of these cases a model immediately precedes the child's own utterance. Imitations being at this age much, much more frequent than spontaneous utterances, these examples are counterbalanced by the single unprovoked [ʏœldə] for *gulden* 'guilder coin', offered by the child when playing with a toy cash register (the *kassa* of (6a)), a form that is the first of its kind (initial fricative in a non-imitation form) in the data. At [1;9–10], more initial fricatives are produced of both types, although still not many.

In (9), (9a) shows all spontaneously produced versions of words containing target fricatives, from [1;3] to [1;8]; (9b) shows all spontaneous output in which an initial fricative is actually produced as such, from [1;8] to [1;10].

(9) Nina: spontaneously produced target fricative data [1;3–9]

- a. 1;3 [tɪtə] zitten 'to sit'
 1;5 [dajt] papegaai 'parrot' [dazət] zuster 'nurse'
 [tas, toəs] bus 'bus' [təutə] veters 'shoe laces'
 [tə] schoen 'shoe' [tat] staat 'stands
 (upright)'
 1;8 [tu(t)s, toets] poes 'pussycat' [kœxə] kuchen 'to cough'
- b. 1;8 gulden 'guilder' [ʏœldə]
 1;9 zaag 'saw': [zax, ʏax] but also: [dax]
 zagen 'to saw': [daxə, dɔʏə, paxə]
 vlinder 'butterfly': [vɪndər, vlɪdər] but also: [lɪndə, dɪnə, tœlə]
 vogel 'bird': [voʏə] but also: [koʏə, okə, oʏo(t),
 tot, ɔx, tɔtə, tots, kɔx]
- 1;10 zo 'so': [zo, ʏo]
 volgende 'next': [vɔ:x]
 foto 'photograph': [foto]

The introduction of onset fricatives makes a slow start, with imitated [fœts] for *fles* at [1;3] antedating spontaneous [ʏœldə] for *gulden* at [1;8] by 5 months. Disregarding the small handful of early imitations with onset-fricatives, the data in (9a), lacking onset-fricatives in early spontaneous output, perfectly reflect the fricative pattern described in the previous section. Clearly, being part of an imitation is not a guarantee for the completely successful production of a sound; yet conversely, for a sound to be part of an imitation may matter as a contribution to its chances of survival. It would be unwise to neglect this factor, especially for the early stages of acquisition considered here.

The issue raised here is not one vigorously debated in recent child speech research reporting. In literature form the early 1970s, however, Edwards (1978: 131ff.), in a thesis on the acquisition of fricatives in English, made a point very similar to that of this section, referring to previous discussions in Ferguson and Farwell (1975) and Garnica and Edwards (1977). The former, in their turn, refer to “discussions of this question from different points of view” in Templin (1947) and Olmsted (1971). The upshot of this literature is as follows. The most principled position in child speech research is to establish a complete separation of imitated from fully spontaneous forms, and to omit the former from consideration. However, Ferguson and Farwell, and – following them – Edwards, point out that especially in studies of very young children few data will survive for analysis when imitations are excluded. The latter decided “to combine spontaneous and imitated forms in order to obtain as complete and accurate a picture of process limitation and suppression as possible. However, it should be mentioned that in different types of studies or in studies with different goals it may still be appropriate (or even necessary) to separate spontaneous and imitated forms” (p.134). The former, whose study covers 3 children in the [0;11] to [1;2] age range, admit that, while “a separation of imitated utterances has not been carried out here, since it would lead to a great reduction of available data without any demonstrable gain of accuracy or homogeneity [...], such a separation might be methodologically sound when dealing with older children, where data are not so limited.” (p.422). Leonard et al. (1978), in a contribution in which, sadly, empirical examples are virtually completely lacking (the research results being presented just in graphs and percentages), compare the spontaneous speech and experimentally solicited nonsense word imitations of 8 English speaking children in the [1;3]–[2;0] age range, and interestingly find two things (p.403): first, children are able to imitate words of a complexity (in syllabic terms) exceeding that of their spontaneous speech, and second “these imitative utterances were subject to the same production constraints”. The suggestion of the current exposition is that the latter claim is close to the truth, but not perfectly so. The remainder of this paper investigates the details and consequences of this latter qualification.

5. Nina’s treatment of *olifant*

One of the analytically most telling forms of the Joost data was his version of the noun *olifant* ‘elephant’ at [2;7], appearing as the output [’ojinɒt]. In Section 2 it was described how this form shows the major components of the Joost analysis in action (onset fricative avoidance, resyllabification, coda adjustment). This item turns out to be of significance in the Nina data, too. It is her renderings of this form that underlie the central points of this paper, so an empirical discussion of this form’s behaviour will be convenient as a start.

The complete collection of Nina data for *olifant* is given in (10) below. These data were assembled from the tape-recorded transcriptions, and are accompanied by notes marking the distinction between direct imitation and more spontaneous speech, assuming this distinction introduced in the previous section to be potentially relevant to the current data. (Also included are a small number of notes from a diary kept simultaneously with the recordings; they are added out of possible interest, but not considered in the discussion because of the lack of information about the precise circumstances in which they were produced). An analysis of Nina's *olifant* output follows the presentation in (10).

- (10) Nina's versions of (adult) *ólifànt* [olifant] 'elephant'
- | | |
|---------------------------------|-----------------|
| a. Taperecorded utterances | b. Diary notes |
| 1;9 1. ɔt | |
| 1;10 2. dæ̃t | |
| 3. fæt [im.] | |
| 4. tæt | |
| 5. bat [im.] | dat |
| 1;11 6. 'oləwə [im.] | |
| (just faintly recorded on tape) | |
| 7. pats | dant |
| 8. 'ofəl̩ant [im.] | fant |
| 2;0 9. tant | |
| 10. f̩ant | 'olf̩antə (PL.) |
| 11. 'owf̩ant | 'oʃ̩f̩ant |
| 2;1 12. 'owf̩ant | 'owəf̩ant |
| 2;2 13. 'hofa | |
| 14. 'owf̩ant | |
| 15. 'olif̩ant | |
| 16. 'olif̩ant | |
| 17. 'olif̩ant | |
| 2;3 18. 'olif̩an-t | 'olif̩ant |

The principal reason for providing this elaborate, and in fact complete, list is that it can be as enlightening to consider in detail the trajectory of a single lexical item in acquisition in progress, as it is to review an overall picture, as was done in the previous two sections. Three stages can be distinguished in these data.

The first stage comprises items 1–5, and ages [1;9–1;10]. The passages containing Nina's utterances at this stage are those in (11).

- (11) Nina: *olifant* at ages [1;9–10]
- [=10.1] Nina, deze? muis.
- at 1;9 Een muis. Weet je deze ook al? ot. [ɔt]
- Een olifant. Een weet je deze? tein.
- [Nina, this one? Mouse. A mouse. Do you know this one, too? [ɔt].
- An elephant. And do you know this one? Train.]

- [=10.2] Weet je die, met een slurf? dut. [dœt]
 at 1;10 Weet je hoe die heet? dut. [dœt]
 Een olifant, ja? Olifant.
 [Do know this one, with a trunk? [dœt]. Do you know what it's called?
 [dœt]. An elephant, yes? An elephant.]
- [=10.3] Dat is een? nipaet. ['nipaet]
 at 1;10 Een olifant. ja. [ja]
 Een olifant, 's geen nijlpaard. baaet. [baet]
 Een olifant fat. [fat].
 [That is a...? ['nipaet]. An elephant. [ja] ('yes'). An elephant, 's not a
 hippopotamus. [baet] An elephant. [fat].]
- [= 10.4/5] En dit dier, ken je dat ook? Weet je hoe die heet? Tat. [tat].
 at 1;10 Een o-li-fant. bat. [bat]
 Niet een tat, een olifant.
 [And this animal, do you know that? What is it called? [tat]. An
 e-le-phant. [bat]. Not a [tat], an elephant.]

This set's first three items show how Nina is apparently unaware of, or at least confused about, the (adult) lexical item *olifant*, at least in her production. In (10.1), [ɔt] is highly likely to represent *hond* [hɔnt] 'dog' and not some radically reduced and modified version of *olifant*; [ɔt] simply is also the child's usual reaction to a picture of a dog, over a dozen times in the files. *Olifant*, however, seems to be a semantically unstable item: elsewhere in the data, this animal gets confused with a hippopotamus (cf. (11[3])), with a monkey, a rhinoceros, and a giraffe.⁷ Consider, as a striking example of such confusion, the contents of the passage leading to output (10[18]):

- (12) [2;3] Nina confusing *olifant* and other animals
 Deze... ε ən ʒɪ'roε
 Nee... ʒɪ'raf
 Nee... n t 'neɪpaet
 Nee, als ze heel groot zijn met een slurf, hoe heten ze dan? ə 'olifan-t
 Goed zo, Nina, ja.
 [This one. Is a giraffe. No. Giraffe. No. A hippopotamus. No, when they are very
 big with a trunk, what are they called then? An elephant. Well done, Nina, yes]

Such instability among all sorts of specific animal names is not uncommon at this stage of development. Smith (1973: 230) notes, for instance, that until [2;6] "horses and donkeys were semantically non-distinct" in his son Amahl's speech, even though two separate words were involved in referring to them, cf. [i:ɔ:] at [2;4] and [gɔki:] at [2;5]. In (11[2]) the output [dœt] represents a different item, in this case *slurf* 'trunk', apparently caused by a misinterpretation of the eliciter's question. In (11[3]) *olifant* is

7. In Fikkert (1994: 161), Jarmo at [2;3] produces for *olifant* the phonologically deviant form [hɔntəflɔŋ]: it contains the same 'dog' element, perhaps by the same sort of confusion.

confused with *nijlpaard* (litt.: Nile horse) ‘hippopotamus’, something that occurs the other way around at [2;2] in a passage incorporating the production of item (10.13). This misunderstanding is resolved by intervention, leading to immediately imitated [fat], which is therefore – at ‘late [1;10]’ – the first *bona fide* occurrence of a version of the item *olifant* in the data.

A comparison, from the data discussed so far, of (10.3)’s imitated [fat] and (10.5)’s imitated [bat] with (10.4)’s apparently spontaneous [tat], fully confirms the suggestion of the previous section that imitations are qualitatively different from spontaneous output in being more faithful to the target: spontaneous speech exhibits the extremely frequent substitute [t-] (see Section 3), whereas the imitations are fully faithful to *f*-, or at least to its labial place of articulation. (For a single item one imitation can obviously be more successful than another, which will not come as a surprise.) This imitation faithfulness can be called premature in the sense that it antedates what happens later in spontaneous speech, cf. [fat] (im.) at [1;10] with [fənt] at [2;0].

When stage 2 in the data of (10) is taken to be items 6–10, ages [1;11–2;0], we can observe a mixture of mono- and polysyllabic forms, with two striking characteristics: the introduction of the final consonant clusters with an increasing degree of accuracy (observe that the diary notes are fully compatible with this developmental fact), and the presence in the data of polysyllabic forms strikingly *only* in imitations. Again, the latter output can be characterised as exhibiting premature faithfulness. Note that item (10.6) [oləwə] is more faithful for *segmental order* than (10.8) [ofələnt]: in the latter, two segments show metathesis, by transposition. Interestingly, the same process in words of a similar shape has been noted earlier in the literature. The forms below for English *elephant* are from Smith (1973), occurring in his son Amahl’s speech (see Smith’s 1973: 36, 98–100, and Appendices A and C):

- (13) *Elephant* in Amahl’s speech (Smith 1973)
 2;3 ɛbininin 2;9 ɛpinin, ɛfələnt, ɛləfənt
 2;4 ɛpinin, ɛbənin(in) 2;10 ɛləfənt
 2;6 ɛpininin

Notice that these Amahl data start exactly at the point in time when the Nina data finish: [2;2] vs. [2;3], which suggests that for the phenomenon at hand, compared to Nina, Amahl seems to have been a slow learner. Although clearly in (13) more is going on, the crucial phenomenon is the early position in the word of the labial in the renderings of *elephant* in the earlier stages. Phonologically perfectly comparable to Nina’s [ofələnt] at [1;11] is Amahl’s [ɛfələnt] at [2;9], showing that Nina is not unique. In all there are 12 occurrences of a variety of forms for *elephant* in Smith’s data; forms occurring once, but undergoing similar treatment, are *aeroplane* [ɛ:bəʔe:n] and *galloping* [gæpəlɪn]. A surprising case of *elephant* metathesis can be found as early as 1926 in A.A. Milne’s classical *Winnie-the-Pooh*. This work’s Chapter Five starts as in (14a), and a passage from later in the same chapter is that in (14b).

(14) *Elephant metathesis in Winnie-the-Pooh*a. CHAPTER FIVE: *In which Piglet meets a Heffalump.*

One day, when Christopher Robin and Winnie-the-Pooh and Piglet were all talking together, Christopher Robin finished the mouthful he was eating and said carelessly: "I saw a Heffalump to-day, Piglet."

- b. Pooh looked around to see that nobody else was listening and said in a very solemn voice: "Piglet, I have decided to do something."
 "What have you decided, Pooh?"
 "I have decided to catch a Heffalump."

Later in the chapter, when Pooh tries to go to sleep by "counting Heffalumps", a picture (by artist E.H. Sheppard) shows an elephant hovering in the air over Pooh in his little bear's bed. It seems that A.A. Milne must have had a very keen ear for child utterances in his surroundings: clearly illustrating metathesis, *heffalump* is Christopher Robin's form for *elephant*.

Finally, Nina's stage 3 comprises all items later than (10[10]), i.e. later than approximately age mid-[2;0]. The output improves from near-faithful to faithful, converging on the correct adult form at the age of mid-[2;2].

The specific topic of the next section will be a more formal analysis of the portion of Nina's grammar that deals with the central phenomena of this and the past few sections. The way in which this will be executed is by a continued case study of the realisations of the item *olifant* (recall (10)), of which there are happily many instances in the files. One of the explicit aims of this exercise will be to return to forms such as [fat] and [oʃələnt] of the current section, imitation forms which were called prematurely faithful in the course of the discussion. Let us proceed and see what this means in terms of a more formal analysis of the child's grammar.


6. A constraint-based analysis of the fricative data

In this section, constraint-based Optimality Theory (henceforth O.T., cf. Prince & Smolensky 1993/2004; McCarthy & Prince 1994; Kager 1999; McCarthy 2002) serves as the theoretical framework of the analysis of the core data of the previous sections. Briefly, in O.T. an analysis, i.e. a partial grammar, takes the form of a hierarchy of constraints. These constraints are universal, but the hierarchy is language-specific. Constraints come in two types: markedness constraints, which express (un)marked properties of a given form (say, of 'a word'), and faithfulness constraints, which limit the distance between an underlying form and its phonetic shape. Typically, the hierarchy of an analysis looks like Phon1 » Faith » Phon2, where Phon1 is the active phonology of the partial grammar in question (in the form of hierarchically ranked universal markedness constraints), and Faith is the hinge point requiring that all phonetic form be identical to underlying form (hence the constraints of Phon2 are the inactive

phonology of the grammar, neutralised by hierarchically earlier Faith). The grammar is used in an evaluation of candidate output forms, and has the task of selecting the correct (optimal) output for the form in question.

As a first analytical step, consider some examples from the relatively stable Joost data, such as the case of output [ni'naf] for target *giráf* [ʒi'raf]. Clearly markedness is involved here, as a nasal is less marked than either a fricative or a liquid: nasals are acquired earlier, for instance, than the other sound classes, especially in onsets (see Fikkert 1998, referring to and largely supporting Beers 1995, for this acquisition sequence in Dutch). In O.T. the analysis of this (relatively simple) case takes the form of the evaluation of candidate output forms by the constraint hierarchy. This is usually visualised in a tableau, as in (15a), which uses the constraints of (15b).

(15) a.

Target (input) /ʒi.'raf/	*OnsCont	SWP	Onset	Faith
1. ʒi.'raf	**			
2. ʒi.'naf	*			*
3. ni.'raf	*			*
4. nir.af			*	*
5.  ni.'naf				**

- b. (i) *ONSCONT: an onset consonant is a plosive or a nasal.
(ii) SWP: stressed syllables are heavy
(iii) ONSET: syllables have onsets
(iv) FAITH(fulness): full identity between input and output.

The tableau includes (i) the underlying form, (ii) the candidate outputs, (iii) the relevant universal constraints (in a left-to-right ranking), and (iv) the evaluation, per candidate, per constraint. The underlying form in this partial child grammar is taken to be the adult surface form, i.e. 'what the child gets' as incoming data. Thus, this underlying form, at least for this stage of acquisition, equals the target. This is a common but not uncontroversial assumption in phonological child speech analysis, but one which will not be disputed here. (For discussion, see for instance Smith 1973: Chapter 4, Vihman 1996: 22ff., and Smolensky 1996).⁸ The candidate outputs are put forward by the Candidate Generator of the grammar. The details of this mechanism will not be discussed here. It will be assumed (counterfactually, but sufficiently correct for present purposes) to provide for evaluation a number of plausible candidate

8. The discussion also ignores the extremely interesting issue of the role of perception in child language, and the influence of perception errors on lexical representations and child production. Again, see the references immediately preceding this footnote in the text, and more recently Pater (2004).

outputs, given the behaviour of the child (for an explanation see Prince and Smolensky 1993/2004: 5ff.). The constraints in the tableau are those listed in (15b), as pointed out above they are provided by Universal Grammar. The remainder of the tableau represents the evaluation procedure of the candidates, and runs as follows. The constraint *ONSCONT assesses candidates 2 and 3 as bad (one *) and candidate 1 as very bad (2 *s), because they have 1 and 2 violations of the constraint, respectively. Candidates 4 and 5 are much better, precisely because they lack fricatives in onsets. SWP, here interpreted as a constraint rather than an input-output mechanism, has no role in the evaluation because in all candidates the stressed syllable is heavy because closed, which is one of the ways in which a syllable can be heavy in languages (in this case in child Dutch), recall Resyllabification in Section 2. ONSET requires output syllables to have onsets (in effect here blocking coda capture by an unstressed syllable). The final stage of faithfulness rejects deviations from the identity between output and input segments. From among the five candidates, [niˈnaf] is clearly worst in terms of faithfulness. But by universal principle the evaluation terminates when just one candidate is left, in this case after the ONSET evaluation stage. The arrow preceding candidate 5 indicates its optimal character: [niˈnaf], although unfaithful to the adult target, is optimal for Joost's grammar. Thus, the avoidance of onset fricatives in this child's phonology follows from the crucial ranking *ONSCONT, onset » FAITH, which is a case of the Phon1 » Faith (» Phon2) design which favours modification of the input.

As a second Joost case, consider how [ˈojinat] can be optimal for *ólifant*, a form whose crucial properties are repeated in (16).

- (16) a. cluster reduction: nt > t (olifant > ojinat) – not analysed here
 b. avoidance of an onset fricative: f > n (olifant > ojinat)
 c. avoidance of a coda liquid: l > j (olifant > ojinat)

Joost avoids liquids in all syllable positions, but strikingly selects different substitutes for liquids in onsets as opposed to codas. Onset liquids are replaced with *n*-, whereas coda liquids are represented by the glide *-j* (see (16b–c)). In [ˈojinat] for *ólifant*, the occurrence of a glide rather than *n*- indicates that the liquid is treated as a coda. Consider the tableau that selects the optimal output for *ólifant*:


(17)

Target /'o.li.fant/	*OnsCont	SWP	Onset	Faith
1. 'o.li.fat	**	*	*	
2. 'o.li.nat	*	*	*	*
3. 'o.ni.fat	*	*	*	*
4. 'o.ni.nat		*	*	**
5. 'o.ji.fat	*		**	*
6. 'oj.if.at			***	*
7. → 'oj.i.nat			**	**

*ONSCONT and SWP make the first selection. The ranking between these two constraints is underdetermined by the data at hand because both accept the optimal output. ONSET and FAITH must be kept out of the way because of their evident dislike of the correct candidate, where ONSET precedes FAITH.

The final Joost case concerns his version of target *útrecht* as ['ytnɛxt]; recall that the naively expected output is *['ytɛxt], given independently active onset cluster reduction.

(18)

Target /'y.tɾɛxt/	*OnsCont	SWP	Onset	Faith
1. 'y.tɾɛxt	*	*	*	
2. 'y.tnɛxt		*	*	*
3. 'yt.rɛxt	*		*	
4.  'yt.nɛxt			*	*
5. 'y.tɛxt		*	*	*
6. 'yt.ɛxt			**	*

After *ONSCONT and SWP, candidates 4 and 6 survive into the second half of the evaluation: there, ['yt.nɛxt] is clearly better for ONSET than ['yt.ɛxt].

This brief analysis of the central Joost data gives a flavour of how O.T. deals with the fricative onset-coda imbalance. The next section contains this paper's case study of the Nina item *ólifant*, with the aim of formalising, and by doing so aiming to gain insight into, her treatment of this form in terms of her (developing) grammar.

7. A constraint-based analysis of the Nina data

Up until early [2;0], Nina's versions of *ólifant* are monosyllabic and fail to licence fricatives in onsets (as do Joost's, *passim*). However, these properties are violated in four forms that were characterised as direct imitations: [fat] and [bat] at [1;10], and ['oləwɑ] and ['ofəlɑnt] at [1;11] (see 10)). Imitation, it was suggested, may lead to premature faithfulness in the data. In addition to developing the relevant portion of Nina's grammar, it is the purpose of this section to show how an O.T. account of the data may reflect this property of her speech.

In O.T. (McCarthy & Prince 1994; Pater 1997) the truncated output characteristic of earlier stages of child speech (such as here monosyllabic versions of *ólifant*) is the result of non-adult (better perhaps: pre-adult) constraint interaction in the grammar. Although usually more constraints are claimed to be involved,⁹ the core interaction is

9. Such as: PARSE-σ: all syllables are parsed into feet; and
FOOT-BIN: all feet are binary (where 'bisyllabic' and 'heavy' may count as binary).

that between *ALIGN*, which favours truncation, and *MAX(IMALITY)*, a branch of faithfulness which rejects deletion. In (19a) are the definitions of these and one further new constraint. Notice how in tableau (19b) the high-ranked position of *ALIGN* ensures output with single feet.

- (19) a. *ALIGN*: A foot edge coincides with a word edge.
MAX: Input segments have corresponding segments in the output.
IDENT(ITY): Segments in the output have identical corresponding segments in the input.
- b. Output: [tat] (Nina truncation stage)

Target /'o.li.fa(n)t/	Align	*OnsFric	Onset	Max	Ident
(o.li)(fat)	**	*	*		
(o.li)(tat)	**		*		*
(fat)		*		*	
☞ (tat)				*	*
(at)			*	*	

In (later Nina and) adult speech, the ranking between *ALIGN* and *MAX* will appear as reversed, preempting truncation. Previous “Faith” has been redefined into separate faithfulness constraints usually assumed in O.T.’s Correspondence Theory (introduced in McCarthy and Prince 1994, also see McCarthy 2002), with *MAX* defined as in (19a), and *IDENT(ITY)* rejecting deviations in candidate segments. In these new terms, the **ONSFRICT* » *IDENT* ranking is common to Joost’s and Nina’s grammars. Referring back to the data list in (10), the ranking of (19b) is proposed to be the grammar underlying Nina’s behaviour up until item (10.9), which – to repeat – has three characteristics: monosyllabic output, no fricatives in onsets, and (prematurely) faithful imitations. The ranking accounts for the first two properties, but clearly fails regarding the third: it explicitly rejects relatively faithful [fat], for instance, which occurs in an imitation. It is proposed here that the latter notion can be incorporated in the following manner: add to the mix the assumption that a direct imitation context may result in temporary promotion of (forms of) Faithfulness. This proposal incorporates two ideas. First, as pointed out, an adult model may but need not enforce a higher degree of Faithfulness: pragmatic or psychological conditions are part of the mix. Moreover, being presented with a model and imitating it may lead to a variety of close approximations to the target, not all of them fully perfect; imitation can be imperfect, hence temporary promotion of Faithfulness can be, too. Second, especially when the output is an imperfect

Further constraints contribute to selecting typically prominent syllables from those of the target word, such as here (a version of) the rightmost prominent (heavy) syllable. See the references just mentioned.

imitation, the *existing grammar* of the child will play a role, rather than that it remains silent while the child makes a more or less successful attempt at parroting. Notice how these ideas are variants of Leonard et al.'s (1977) intuition, mentioned at the end of Section 4, that the “imitative utterances [of children a]re subject to the same production constraints” as their spontaneous speech is – but now developed within a theoretical framework. Tableau (20) visualises how promotion of Faithfulness can account for (imperfect) imitation producing prematurely faithful [fat] in Nina's monosyllabic stage, cf. (10[3]):

(20) Output: [fat] (imitation tableau) [1;10]

Target /'o.li.fa(n)t/	(Faith-)Ident	Align	*OnsFric	(Faith-)Max
(o.li)(fat)		*	*	
(o.li)(tat)	*	*		
☞ (fat)			*	*
(tat)	*			*

The imitation is expressed as promotion of IDENT, in a grammar with the ranking of (19b); as a result, the ranking reversal of *ONSFric and IDENT produces an output more faithful than would otherwise be the case for the grammar. The imitation's imperfection is represented by the preservation of ALIGN » MAX.

A number of examples occurring in imitation contexts¹⁰ are more perfect, suggesting that IDENT and MAX can be simultaneously promoted. A case in point is [1;5] [pəpxaj] for *papegaai* ‘parrot’ (recall Section 4), which is virtually completely faithful in number and shape of segments. Another case is [fœts] for *fles* ‘bottle’ at [1;3], which violates not a single constraint of (20), although it occurs at a stage at which *ONSFric is still highly active. Finally, [’oləwə] for *olifant* at [1;11], a polysyllabic imitation at what is still the monosyllabic stage, may be considered completely faithful: this utterance's latter half is extremely faint in the recording, implying that the transcribed form is incomplete rather than unfaithful in a grammatical sense.

In all variants discussed in these sections, for both Joost and Nina, the early child grammar is essentially a substitution grammar with crucial *ONSETFric » Faith, in which typically the offending continuant is replaced in situ with a non-continuant. For Joost substitution is routine, for Nina it is the preferred output adjustment by far. Recall from the reference list in section 1, however, that other ways are available of avoiding onset fricatives: one of these is metathesis, in (English) forms with transposed segments such as [lɪf] for *story* mentioned in [1], and [uz] for *zoo* in [7]. As shown

10. The imitation form [bat] will not be separately analysed, it will be assumed to be equivalent of [fat] in all currently relevant respects (i.e., involving promotion of more fine-grained featural IDENT-LABIAL, for instance).

above, Nina in fact employs metathesis at a slightly later stage, in premature faithfulness in an imitation. It will now be shown how the current approach is able to capture this behaviour.

O.T. literature (McCarthy & Prince 1994; McCarthy 1995; 2000, Hume 1998; Dos Santos and Rose 2006) suggests that cases of metathesis be handled as violations of a particular type of faithfulness, namely the linear faithfulness holding in the left-to-right arrangement of segments, comparing input and output. The proposal separates out, from among the variety of types of Faithfulness, a constraint LINEARITY (see (21)), which is usually high-ranked among the faithfulness constraints, metathesis being uncommon in the languages of the world; but, simply by being a constraint, it can also be low(er)-ranked and violated by optimal output. Metathesis occurs in Nina's output ['ofəlɒnt], but in order to simplify the exposition, capturing metathesis in O.T. will first be illustrated using the simpler English form *zoo* surfacing as [uz], for which transposition is motivated by a desire to avoid fricatives in onsets, and hence the ranking *ONSFric » IDENT is again the point of departure of the exercise. Given these assumptions, there are three possible rankings for LINEARITY, cf. (22).

- (21) LINEAR(ITY): The precedence structure of the output is consistent with that of the input, and vice versa.

- (22) a. LINEAR » *ONSFric » IDENT (substitution)

Target: /zoo/	(Faith-)Linear	*OnsFric	(Faith-)Ident
Zoo		*	
Ooz	*		
ʔ too			*
Oot	*		*

- b. *ONSFric » LINEAR » IDENT (substitution)

Target: /zoo/	*OnsFric	Linear	Ident
Zoo	*		
Ooz		*	
ʔ too			*
Oot		*	*


- c. *ONSETFric » IDENT » LINEAR (metathesis)

Target: /zoo/	*OnsFric	Ident	Linear
Zoo	*		
ʔ ooz			*
Too		*	
oot		*	*

Comparing the rankings of (22), a substitution grammar is characterised not only by *ONSFric » IDENT, but also by LINEAR » IDENT, with the ranking of *ONSFric and LINEAR undecided. In other words, there are two possible ways of formalising a substitution grammar: as either (22a) or as (22b). Tableau (22c) represents a grammar which consistently solves avoiding fricatives by metathesis. Its crucial property is IDENT » LINEAR, which illustrates the common O.T. practice of eliminating the competition in order to arrive at the optimal candidate.

Assuming that the bulk of her output shows that basically Nina has a substitution grammar, and continuing to assume that the way to account for imitation is by promoting Faith to a higher-ranked position, let us consider what Nina's metathesis output ['ofəlant] tells us about her grammar: it occurs in imitation circumstances, but is otherwise completely atypical of her speech. Let us reiterate this form's relevant properties. First, a fricative fails to occur in its target onset position, suggesting the relevance of *ONSFric. Second, the fricative has found a landing site familiar from the first half of this paper, namely immediately to the right of a stressed vowel; this suggests the relevance of SWP. Third, although the form is produced right in the middle of what appears to be the monosyllabic stage, it is perfectly faithful with respect to both segments and number of syllables; this indicates high-ranked Faithfulness, with the exception of LINEAR. In sum, it is just the metathesis that makes this form odd. Reviewing the two substitution grammars of (22a)–(22b), it is the latter and only the latter which – interestingly enough – can accommodate metathesis by promotion of Faith: IDENT top-ranked selects ooʒ; the same operation performed on (22a) simply selects faithful zoo. Gratefully accepting this result, the tableau in (23) gives the analysis of Nina's output:

(23) ['ofəlant] (imitation metathesis tableau for monosyllabic stage) at [1;11]

Target /'o.li.fant/	Max/Ident	Align	*OnsFric	Onset	Linear
1. 'ol.ə.fant		*	*	**	
2.  'of.ə.lant		*		**	*
3. 'ol.ə.tanf		*		**	*
4. 'o.fə.lant		*	*	*	*
5. 'o.lə.tant	*	*		*	
6. monosyllabic	*				

The promoted MAX/IDENT pair ensures complete segmental faithfulness, preempting ALIGN as the source of monosyllabic output (the further potential properties of which vis-à-vis the constraints are therefore irrelevant and not represented here, see the empty cells for monosyllabic candidate family 6). *ONSFric » LINEAR enforces movement of the fricative. The result is a metathesised form as optimal, as required. Notice two things, however. First, the tableau reveals the analysis to be incomplete, in

that it does not explain why, in this long word with more than one available coda, the fricative moves to the left rather than to the right – the result being that candidate 3 is just as optimal as candidate 2. Second, this proposal relies on substitution grammar (22b) which allows metathesis, a property (22a) lacks; this makes one curious as to the source of this apparently preferred grammar, which is unlikely to come just from nowhere. These two points, and (sketches of) approaches to their solution, will be addressed in the remainder of this paper, in the order given.

In O.T. the usual reaction to an incomplete analysis such as that in (23) is to search for a structural difference between the two candidates as a result of which one can be called more optimal than the other, and to suspect the interference of an independent constraint which targets this difference. Two possibilities spring to mind, one involving the right edge of the input, the other the left edge. First, consider the right edge. A crucial aspect of suboptimal candidate 3 is that it disturbs the alignment of the right edge and the final segment of the word: edges are often special in preventing phonological processes from doing this damage. In O.T. this condition has been formulated as *ANCHOR(ING)*, which demands the underlying edge segment to reoccur in the output in the same position (see, for instance, McCarthy & Prince 1994, from which this notion originates). *ANCHOR* added to (23), bottom-ranked following *LINEAR* for instance, excludes metathesis towards either edge, eliminating candidate 3 (and vacuously eliminating any fricative-initial candidate, given **ONSFRICT*). This leaves word-internal metathesis as the optimal way of complying with **ONSFRICT*, assuming low-ranked *LINEAR*. This proposal enjoys some independent motivation. *ANCHOR* does not care about a candidate's size, it just addresses edges: thus, it also provides a means of counteracting in a grammar of type (22b) absent edge-affecting metathesis in truncated monosyllabic output: *[oli-]/fant/* → **[taf]*. The analysis eventually combines linear faithfulness in monosyllabic imitation output with metathesis in polysyllabic imitation output.

The left edge view of a form such as *[ʰofəlant]* resorts to moving the fricative for a reason different from **ONSFRICT*. There are in fact two independent incentives for finding such an alternative. First, as pointed out in Section 3, the Nina data only weakly support the word-internal operation of **ONSFRICT*. Second, at least some O.T. literature proposes that this latter constraint not stand on its own, but be subsumed in a set of so-called fixedly ranked constraints regulating preferred onset consonants, using a version of the sonority hierarchy (Gnanadesikan 1995/2004; Pater 1997; Pater & Barlow 2003). Although details of the proposals differ, this ranked set can be taken to have the following form:

- (24) Preferred onsets: **ons-gl* » **ons-liq* » **ons-nas* » **ons-fric* » **ons-plos*

This fixedly-ranked set formalises the claim that universally liquids are worse onsets than fricatives, with very ill consequences for the **ONSFRICT* explanation of *[ʰofəlant]* metathesis: the second member of the output transposition pair is a liquid, an onset

even worse than the dispreferred fricative.¹¹ (By the same token, the substitution consonant is highly preferred to be a plosive and not any of the sonorants, underscoring the oddity of Joost's choice in this respect.) Reviewing the literature, another candidate constraint able to move /f/ is one proposed by Velleman (1996: 176–177) in a “feature-by-position” discussion trying to account for the relatively old observation (Ingram 1974; Macken 1979) that many “children produce consonant place features in a[n articulatorily] front-to-back order within a word” (Ingram 1974: 149; also see Ingram 1978 and Macken 1979). One of the ways in which this effect shows up happens to be metathesis, as in (some children's) English [mɪn] for *cream*, [piti] for *t.v.*, and [baki] for *coffee*. Assuming a constraint LABIAL-LEFT, in the current case it will be /f/'s place of articulation rather than its manner that causes its displacement. ANCHOR will again be the constraint prohibiting the labial from reaching the word edge. Some small pieces of support for LABIAL-LEFT in Nina's grammar come from data such as [ˈdœfələ] for *dolfijn* [dɔlˈfeɪn] ‘dolphin’ at [1;9], and [ˈɔfəja, ˈowəvar] for *óievaar* [ˈoɣəvar] ‘stork’ at [1;10].¹² A more sophisticated and deeper account of the possible ways of formulating Nina's metathesis trigger will be left for future research. The currently more important conclusion is that there seems to be one grammar, namely (22b), which can be claimed to underlie her behaviour, and which can be shown to successfully cover a variety of cases of imitation faithfulness, by promoting Faith. Thus, even imitation is a phenomenon under grammatical control.

8. Epilogue

The analysis of the previous sections, however sketchy and incomplete, has an intriguing consequence which is worth winding up this paper with. Nina's early grammar (22b) has the ranking *ONSFRICT » LINEAR » IDENT. This ranking is compatible with consistent substitution output, and with two sorts of faithful imitation output (accounting for them by temporary high-ranking of (forms of) Faith): a truncated but otherwise highly faithful form if ALIGN is left unviolated at [1;10], and a polysyllabic metathesised form violating ALIGN at [1;11]. ANCHOR is assumed to prohibit edge violation throughout. This grammar, which was selected on the basis of an interaction between a variety of empirical data and a number of assumptions on how to deal with them,

11. Do recall, though, a metathesised example such as [ɔɪs] for *story* mentioned in [1] of the survey in the Introduction; from the point of view of such data the sonority hierarchy, and the position of individual sounds in it, seems worth further looking into.

12. LABIAL-LEFT may also be the explanation of the *olifant* form [ɔfəsant] mentioned by Fikkert (1994: 227) for Tom at [1;7], in which *ONSFRICT is violated by a fricative placeholder.

has an interesting property which (22a) lacks: rather than its rival, it is compatible with an O.T. assumption about the essential properties of early first language acquisition, to the effect that the family of ‘markedness constraints’, to which *ONSFRICT (and LABIAL-LEFT, perhaps) belongs, start out dominating the family of faithfulness constraints, which incorporates LINEAR, MAX, and IDENT. Reasons for making this assumption can be found laid out in literature such as Gnanadesikan (1995/2004), and Tesar and Smolensky (1998), which points out that this architecture explains the generally ‘unmarked’, as opposed to faithful, character of early child speech. The currently important point is that no language-specific learning is involved in the core fragment of the analysis: the ranking *ONSFRICT » (different forms of) Faith comes for free in the grammar. In addition to this, among the faithfulness constraints, grammar (22b) requires LINEAR » IDENT.¹³ This is a faithfulness-internal ranking which must have some source as well. Speculating, observe that assigning internal architecture to the set of faithfulness constraints is not unheard of: Smith (2000) shows that in the O.T. variant espousing so-called “Positional Faithfulness” (Lombardi 1999; Alderete 2003), the acquisition starting point must be PosFaith » Faith. It is not unimaginable that LINEAR » IDENT is a similar case of initial ranking.

In this paper, idea such as those described in this final section received support from a number of central assumptions, to wit: first, the notion that the distinction between spontaneous and imitation output is relevant to phonological analysis, and second, that the latter, seemingly cognitive, process can and should be formalised in terms of the grammar, as temporary (and typically premature) promotion of Faith. In Nina’s grammar, the ongoing acquisition process reranks constraints in order to conform more adequately to its many and rapidly increasing number of targets. A constraint such as *ONSFRICT will be assigned a position below Faith, and eventually veritably spontaneous output will become indistinguishable from imitations.

References

- Alderete, J.D. 2003. Structural disparities in Navajo word domains: A case for LEXCAT-FAITHFULNESS. *The Linguistic Review* 20: 111–158.
- Arias, J. & Lleó, C. 2007. On the unmarkedness of Stress-to-Weight. System und Variation Workshop, Annual Meeting of the deutsches Gesellschaft für Sprachwissenschaft 2007, Universität Siegen.
- Beers, M. 1995. The Phonology of Normally Developing and Language-impaired Children. Ph.D. dissertation, University of Amsterdam.

13. I am grateful to a referee for urging the author to dwell on this component of the analysis, too.

- Bernhardt, B.H. (ed). 2002. Special issue: Intervocalic consonants. *Clinical Linguistics and Phonetics* 16: 149–232.
- Bernhardt, B.H. & Stemberger, J.P. 1998. *Handbook of Phonological Development. From the Perspective of Constraint-based Nonlinear Phonology*. San Diego CA: Academic Press.
- Chiat, S. 1989. The relation between prosodic structure, syllabification, and segment realization: Evidence from a child with fricative stopping. *Clinical Linguistics and Phonetics* 3: 223–242.
- Crosswhite, K. 1998. Segments vs. prosodic correspondence in Chamorro. *Phonology* 15: 281–316.
- Dinnsen, D.A. 1996. Context effects in the acquisition of fricatives. In *Proceedings of the UBC International Conference on Phonological Acquisition*, B. Bernhardt et al. (eds), 136–148. Somerville MA: Cascadilla.
- Dos Santos, C. & Rose, Y. 2006. Influences multiples dans l'harmonie consonantique et la métathèse en acquisition du français. *Recherches Linguistiques de Vincennes* 35: 77–102.
- Edwards, M.L. 1978. Patterns and Processes in Fricative Acquisition: Longitudinal Evidence from six English-learning Children. Ph.D. dissertation, Stanford University.
- Edwards, M.L. 1996. Word position effects in the production of fricatives. In *Proceedings of the UBC International Conference on Phonological Acquisition*, B. Bernhardt et al. (eds), 49–158. Somerville MA: Cascadilla.
- Farwell, C.B. 1976. Some strategies in the early production of fricatives. *Papers and Reports on Child Language Development* 12: 97–104.
- Ferguson, C.A. 1978. Fricatives in child language acquisition. In *Papers on Linguistics and Child Language*, V. Honsa & J. Hardman-de-Bautista (eds), 93–115. The Hague: Mouton.
- Ferguson, C.A. & Farwell, C.B. 1975. Words and sounds in early language acquisition. *Language* 51: 419–439.
- Ferguson, C.A., Peizer, C.B. & Weeks, T.E. 1973. Model-and-replica phonological grammar of a child's first words. *Lingua* 31: 35–65.
- Fikkert, P. 1994. On the Acquisition of Prosodic Structure. Ph.D. dissertation, Leiden University.
- Fikkert, P. 1998. The acquisition of Dutch phonology. In *The Acquisition of Dutch*, S. Gillis & A. de Houwer (eds), 163–222. Amsterdam: John Benjamins.
- Garnica, O.K. & Edwards, M.L. 1977. Phonological variation in children's speech: the trade-off phenomenon. *Ohio State University Working Papers in Linguistics* 22: 81–87.
- Gierut, J.A. 1985. On the Relationship between Phonological Knowledge and Generalization Learning in Misarticulating Children. Ph.D. dissertation, Indiana University. (Reproduced by Indiana University Linguistics Club, 1986).
- Gnanadesikan, A. 1995 [2004]. Markedness and faithfulness constraints in child phonology. Ms, Rutgers University. Revised version published in 2004, *Constraints in Phonological Acquisition*, R. Kager, J. Pater & W. Zonneveld (eds), 73–108. Cambridge: CUP.
- Goldsmith, J.A. 1990. *Autosegmental & Metrical Phonology*. Oxford: Blackwell.
- Grijzenhout, J. & Joppen, S. 1998. Syllable structure and the acquisition of German consonants: A case study. Ms, Heinrich-Heine-Universität, Düsseldorf.
- Hammond, M. 1992. Resyllabification in English, Ms, University of Arizona, Tucson.
- Hammond, M. 1999. *The Phonology of English: A Prosodic Optimality-Theoretic Approach*. Oxford: OUP.
- Hammond, M. & Dupoux, E. 1996. Psychophonology. In *Current Trends in Phonology: Models and Methods*, J. Durand & B. Laks (eds), 281–304. Salford: University of Salford.
- Hume, E. 1998. Metathesis in phonological theory: The case of Leti. *Lingua* 104: 147–186.
- Ingram, D. 1974. Fronting in child phonology. *Journal of Child Language* 1: 233–241.

- Ingram, D. 1978. The role of the syllable in phonological development. In *Syllables and Segments*, A. Bell & J.B. Hooper (eds), 143–155. Amsterdam: North-Holland.
- Ingram, D. 1989. *First Language Acquisition. Method, Description and Explanation*. Cambridge: CUP.
- Kager, R. 1999. *Optimality Theory*. Cambridge: CUP.
- Leonard, L.B. & McGregor, K.K. 1991. Unusual phonological patterns and their underlying representations: A case study. *Journal of Child Language* 18: 261–272.
- Leonard, L.B., Schwartz, R.G., Folger, M.K. & Wilcox, M.J. 1978. Some aspects of child phonology in imitative and spontaneous speech. *Journal of Child Language* 5: 403–415.
- Leopold, W. 1947. *Speech Development of a Bilingual Child: A Linguist's Record*. Evanston IL: Northwestern University Press.
- Lohuis-Weber, H. & Zonneveld, W. 1996. Phonological acquisition and Dutch word prosody. *Language Acquisition* 5: 245–284.
- Lombardi, L. 1999. Positional faithfulness and the phonology of voicing in optimality theory. *Natural Language and Linguistic Theory* 17: 267–302.
- McCarthy, J.J. 1995/2000. Extensions of faithfulness: Rotuman revisited. Ms, University of Massachusetts, Amherst. Revised version published in 2000, The prosody of phase in Rotuman. *Natural Language and Linguistic Theory* 18: 147–197.
- McCarthy, J.J. 2002. *A Thematic Guide to Optimality Theory*. Cambridge: CUP.
- McCarthy, J.J. & Prince, A.S. 1994. The emergence of the unmarked: Optimality in prosodic morphology. In *Proceedings of NELS 24*, M. Gonzalez (ed), 333–379. Amherst MA: University of Massachusetts.
- Macken, M.A. 1979. Developmental reorganization of phonology: A hierarchy of basic units of acquisition. *Lingua* 49: 11–49.
- Marshall, C. & Chiat, S. 2003. A foot domain account of prosodically conditioned substitutions. *Clinical Linguistics and Phonetics* 17: 645–657.
- Menn, L. 1971. Phonotactic rules in beginning speech. *Lingua* 24: 225–251.
- Milne, A.A. 1926 [1977]. *Winnie-the-Pooh*. Illustrated by E.H. Shepard. 13th reprint of paperback edn. London: Methuen.
- Myers, S. 1987. Vowel shortening in English. *Natural Language and Linguistic Theory* 5: 485–518.
- Olmsted, D.L. 1971. *Out of the Mouth of Babies*. The Hague: Mouton.
- Pas, B. van der. 1997. Fricative Avoidance in Phonological Acquisition. MA thesis, Utrecht University.
- Pater, J. 1994. Against the underlying specification of an 'exceptional' English stress pattern. Ms, McGill University.
- Pater, J. 1997. Minimal violation and phonological development. *Language Acquisition* 6: 201–253.
- Pater, J. 2004. Bridging the gap between receptive and productive development with minimally violable constraints. In *Constraints in Phonological Acquisition*, R. Kager, J. Pater & W. Zonneveld (eds), 219–244. Cambridge: CUP.
- Pater, J. & Barlow, J.A. 2003. Constraint conflict in cluster reduction. *Journal of Child Language* 30: 487–526.
- Pearce, M. 2006. The interaction between the metrical structure and tone in Kera. *Phonology* 23: 259–286.
- Prince, A.S. 1990. Quantitative consequences of rhythmic organization. In *Papers from the 26th Regional Meeting of the Chicago Linguistic Society* [The Parasession on the Syllable in Phonetics and Phonology2], M. Ziolkowski et. al. (eds), 355–398. Chicago IL: CLS.

- Prince, A.S. & Smolensky, P. 1993 [2004]. *Optimality Theory: Constraint Interaction in Generative Grammar*. New Brunswick NJ: Rutgers University, New Brunswick & Boulder CO: University of Colorado. Published in 2004, *Optimality Theory, Constraint Interaction in Generative Grammar*. Oxford: Blackwell.
- Selkirk, E.O. 1982. The syllable. In *The Structure of Phonological Representations*, Part II, H. van der Hulst & N. Smith (eds), 337–384. Dordrecht: Foris.
- Smith, J. 2000. Positional Faithfulness and learnability in optimality theory. In *Proceedings from the Eastern States Conference on Linguistics*, R. Daly & A. Riehl (eds), 203–214. Ithaca NY: CLC.
- Smith, N.V. 1973. *The Acquisition of Phonology, a Case Study*. Cambridge: CUP.
- Smolensky, P. 1996. On the comprehension/production dilemma in child language. *Linguistic Inquiry* 27: 720–731.
- Stemberger, J.P. 1996. Syllable structure in English, with emphasis on codas. In *Proceedings of the UBC International Conference on Phonological Acquisition*, B.H. Bernhardt et al. (eds), 62–75. Somerville MA: Cascadilla.
- Stoel-Gammon, C. 1985. Phonetic inventories, 15–24 months: A longitudinal study. *Journal of Speech and Hearing Research* 28: 505–512.
- Stoel-Gammon, C. & Cooper, J.A. 1984. Patterns of early lexical and phonological development. *Journal of Child Language* 11: 247–271.
- Templin, M.C. 1947. Spontaneous versus imitated verbalization in testing articulation in pre-school children. *Journal of Speech Disorders* 12: 293–300.
- Tesar, B., & Smolensky, P. 1998. Learnability in optimality theory. *Linguistic Inquiry* 29: 229–268.
- Velleman, S.L. 1996. Metathesis highlights feature-by-position constraints. In *Proceedings of the UBC International Conference on Phonological Acquisition*, B. Bernhardt et al. (eds), 173–186. Somerville MA: Cascadilla.
- Vihman, M.M. 1996. *Phonological Development*. Oxford: Blackwell.
- Waterson, N. 1971. Child phonology: A prosodic view. *Journal of Linguistics* 7: 179–221.
- Zonneveld, W. 1999. Heffalumps: Onset fricatives in child speech. Ms, UiL-OTS, Utrecht University.
- Zonneveld, W., van der Pas, B. & de Bree, E. 2007. Feet and syllables in elephants and missiles. *Clinical Linguistics and Phonetics* 21: 29–36.

Index

A

- A-bar chain 115–138, *see also*
chain formation
- acquisition 1–10, 66, 74,
115–129, 133, 136, 138,
204–206, 327–338, 350,
358–360, 364, 373
- of fricatives 347, 359
- phonological 345–373
- second language 8, 330–331,
334, 336
- Acquisition by Processing
Theory (APT) 338
- agreement 88–90, 99, 102,
116–123, 137–138, 236,
258–260, 265–271, 284,
315–318, 322–325, 332
- relative 116, 120, 122,
127–129, 131–133
- anaphora 6, 8, 49, 51, 141,
144–145, 148, 160–163, 243,
247–250, 297
- complement 243, 247–250
- anaphoric dependency 143,
145, 152, 154, 162

B

- Binding Theory 142, 144, 189,
261, 278
- bound variable 257–258,
260–262, 269–270,
273–279, *see also* variable
binding

C

- case marking 179, 316
- case
 - morphological 124
 - oblique 117, 121–122, 132
 - structural 80, 82, 84–86,
94–96, 117, 121–122,
265, 277
- causal relation 168, 173, 176,
178–181

Categorical Grammar

- see* Symmetric Categorical
Grammar
- chain formation 143, 145–146,
162–164, 257, 263–266,
268, 271, 276–277, 279,
see also A-bar chain
- child speech 345–348, 350–351,
359, 364, 366, 373
- clitic 81, 87, 267–268, 274,
315–325
- competence 6, 8–10, 183–189,
196, 205, 208–209, 295,
327–331, 336
- complementizer 116,
119–120, 127
- computation 51, 111, 184,
186–187, 195–196, 204,
269–270, 275, 291–297, 304
- computational efficiency
183, 195
- connective
 - causal 174–180
 - temporal 175–176, 179
- connectivity 288
- constraint ranking and
reranking 345–373
- context 1, 7, 12–14, 34, 50–51,
58, 60, 67–68, 89, 92,
100–113, 136, 159, 176,
189–194, 203, 209,
229–232, 238–249, 260,
265–266, 279, 287–299,
319–321, 332–333, 342–347,
357, 367
- continuation semantics 41,
49, 50–52
- control 67, 72, 74, 88–95,
146–151, 229–231, 265,
341, 372
- coreference 146, 149–152,
154–157, 161–164, 257,
261–262, 271, 275, 277
- cyclicity 115–138

D

- directivity 241–244, 246–247,
249–250
- disambiguation 170, 172–174,
177–178
- discourse 3–8, 113, 141–150,
156, 160–168, 180–181,
208–209, 242–243, 261,
270–275, 297, 307–308,
319–324, 337
- Dutch 5, 58–63, 74, 79–83,
86, 91–95, 104–138, 144,
147, 151–162, 173–174, 213,
219–227, 235, 263–269,
278–279, 293–294,
305–307, 324, 345–348, 353,
364–365

E

- economy 6, 10, 125, 141–151,
155–156, 160–161, 257–258,
262, 268–276, 279
- ellipsis 103, 156, 158–159, 273,
278–279
- English 5–6, 18, 58–65, 71–74,
83, 105, 108–119, 133, 144,
147, 151, 156, 162, 173–174,
222, 235, 260, 266–279,
335, 340, 346–350, 355, 359,
362, 368–372
- Canadian 59
- extended projection 79, 82,
88–92, 95–96
- eye-tracking 4, 6, 141–142,
146–154, 157–162, 273, 275

F

- faithfulness 345–348, 351, 357,
362–373
- feature 39, 57, 60–65, 68–71,
126, 129–135, 145, 188–194,
257, 260–268, 276–279,
338, 372
- animate 116–133, 316–317

- gender 317, 319–325
 number 266, 319–320, 323
 person 258, 260, 265
 fixed expression 295, 297, 301, 305, 308
 French 79, 116, 316–318, 324–325
- G**
 Generalized Quantifier 244–247
 generator 184–189, 196, 201–208, 364
 German 80, 83–88, 94–95, 119–120, 124, 134, 138, 257–260, 263, 268, 273, 279, 346–347
 government transparency 82, 87
 grammar 1–6, 10–18, 24, 49, 79–83, 88, 96–97, 104–106, 115, 123, 133, 138, 183–196, 201–209, 265–274, 291–296, 306, 327–342, 345–348, 363–373
 compilation and interleaving 183–120
- H**
 head movement 83, 87, 91, 94
- I**
 idealization 9
 idiom processing 291–308
 imitation 345–346, 356–363, 366–373
 inclusiveness 189–196, 201–202, 204
 infant word learning 55–75
 Infinitivus pro Participio 95
 interface 16–17, 30, 32, 41, 51, 257–261, 270, 278–279, 329–330, 337, 340
 island 135–137
 isomorphism 101, 105, 108
 Italian 316, 318, 325
- L**
 language representation 291–308
 Left Branch Condition 119, 123–125, 127
 left edge 94, 125–126, 136–137, 316, 371
- level of description 185–188, 193, 196, 200–201, 204, 206
 lexical bias 213–215, 222
 linearization options 80
 locality 87, 90, 94–96, 115, 120, 137–138, 143, 257–260, 263, 269, 278–279
 logophors 143, 146–147, 161–162, 164
 long passives 83, 93, 96
- M**
 mapping 18, 99, 113, 185–186, 197, 199, 201, 204, 238
 syntax-semantics mapping 18, 99, 113
 metalanguage 21, 327, 330, 334
 metalinguistic ability 327, 330, 341
 mirror effect 80, 94
 modified numerals 248, 251–253
 Modular On-line Growth and Use of Language (MOGUL) 327–328, 336, 338, 341–342
 monitoring 207, 213–217, 221–223, 228–232, 271
 monotonicity 27, 246–250, 254
- N**
 noticing hypothesis 332
- O**
 object
 direct 25, 93, 315, 318–325
 indirect 317, 324
 one-way weak reciprocity 288
- P**
 parser 12, 95–96, 99, 106, 112–113, 145, 154–159, 162, 185–196, 201, 204–208, 329, 339
 partitioned interpretation 283–289
 past participle 94–96
 performance 6–10, 60–67, 71–72, 138, 156, 183–188, 201–209, 215, 273, 296, 300–301, 327–335, 341–342
 phonetic detail 55–56, 67–74
- pied-piping 91–93, 115, 119, 124–127, 136–137
 plurality 266, 284, 286, 289
 Portuguese 267, 274, 316
 preferential looking paradigm 55–75
 PRO 89–91, 95, 126, 265, 350
 processing 5, 8, 66–72, 141, 144–157, 159–169, 173–185, 199, 201, 208, 271–274, 279, 291–301, 307–308, 327–342
 pronoun 27, 115–122, 125–137, 142–144, 149–155, 162, 169–180, 243, 258–262, 266, 269–271, 297, 315–322, 333
 A-bar 116–117, 121, 126, 128–129, 131–132, 136–137
 intermediate A-bar 129, 132
 personal 170–172, 177
 relative 27, 119, 121, 129, 131
 proper name conjunctions 283, 287, 289
 pruning 87
 psychological reality 6–83, 97, 121, 138, 141–142, 162–164, 183, 214, 315, 328–329
- Q**
 quantification 16, 156, 238, 251, 254
 quantifier 3, 7, 11–13, 16, 26–32, 37–42, 45–51, 102, 108, 152–154, 235–254, 283–289, 321–323
 Question Under Discussion 7, 99, 106–109, 111–113
- R**
 reading
 strict 156–157, 159, 258, 261, 273
 sloppy 157, 159, 258
 reciprocal expression 283–288
 reflexives 141–143, 146–149, 161–164, 266, 268
 relatives 120, 122, 130, 133, 137, 203
 restructuring
 argument 80
 head 82

Rule I 257, 260–262, 268–271,
273, 279

S

scope

inverse 39, 41–44, 52,
99–111, 113

surface 7, 99–111, 113

Spanish 5, 116, 315–321,
324–325, 350

speech error 215, 222,
225, 230

speech

inner 213–219, 221–232

overt 213–219, 221–225,
227–232

spontaneous 213, 218–223,
226–232, 335, 341, 345, 348,
356–360, 362, 368

storage 55, 149, 291–296, 304

stranding 92, 125, 136–137

Stress-to-Weight

Principle 349–351

Strongest Meaning

Hypothesis 284

structure building 82, 188–189,
192, 194

structure destroying 82

sub-extraction 81

switch paradigm 56, 58, 61,
63, 71

syllable weight 345–373

Symmetric Categorical

Grammar 11

T

tacit knowledge 5, 330,
333–334, 336–338

terms and contexts 51

text processing 291–308

V

vagueness 240–241, 340

variable binding 146,
149–151, 155–157, 162–164,
257, 261, 269–272, 275,
277, *see also* bound
variable

verb cluster 89–97

W

wh-movement 88, 95, 115, 120,
123, 126–127, 133–138

wh-question 115–116, 136

Z

zu/te 83, 90–91

Linguistik Aktuell/Linguistics Today

A complete list of titles in this series can be found on the publishers' website, www.benjamins.com

- 152 **GALLEGO, Ángel J.:** Phase Theory. xii, 360 pp. + index. *Expected February 2010*
- 151 **SUDHOFF, Stefan:** Focus Particles in German. Syntax, prosody, and information structure. xiii, 323 pp. + index. *Expected February 2010*
- 150 **EVERAERT, Martin, Tom LENTZ, Hannah De MULDER, Øystein NILSEN and Arjen ZONDERVAN (eds.):** The Linguistics Enterprise. From knowledge of language to knowledge in linguistics. 2010. ix, 379 pp.
- 149 **AELBRECHT, Lobke:** The Syntactic Licensing of Ellipsis. xii, 229 pp. *Expected January 2010*
- 148 **HOGEWEG, Lotte, Helen de HOOP and Andrej MALCHUKOV (eds.):** Cross-linguistic Semantics of Tense, Aspect, and Modality. 2009. vii, 406 pp.
- 147 **GHOMESHI, Jila, Ileana PAUL and Martina WILTSCHKO (eds.):** Determiners. Universals and variation. 2009. vii, 247 pp.
- 146 **GELDEREN, Elly van (ed.):** Cyclical Change. 2009. viii, 329 pp.
- 145 **WESTERGAARD, Marit:** The Acquisition of Word Order. Micro-cues, information structure, and economy. 2009. xii, 245 pp.
- 144 **PUTNAM, Michael T. (ed.):** Towards a Derivational Syntax. Survive-minimalism. 2009. x, 269 pp.
- 143 **ROTHMAYR, Antonia:** The Structure of Stative Verbs. 2009. xv, 216 pp.
- 142 **NUNES, Jairo (ed.):** Minimalist Essays on Brazilian Portuguese Syntax. 2009. vi, 243 pp.
- 141 **ALEXIADOU, Artemis, Jorge HANKAMER, Thomas McFADDEN, Justin NUGER and Florian SCHÄFER (eds.):** Advances in Comparative Germanic Syntax. 2009. xv, 395 pp.
- 140 **ROEHRS, Dorian:** Demonstratives and Definite Articles as Nominal Auxiliaries. 2009. xii, 196 pp.
- 139 **HICKS, Glyn:** The Derivation of Anaphoric Relations. 2009. xii, 309 pp.
- 138 **SIDDIQI, Daniel:** Syntax within the Word. Economy, allomorphy, and argument selection in Distributed Morphology. 2009. xii, 138 pp.
- 137 **PFAU, Roland:** Grammar as Processor. A Distributed Morphology account of spontaneous speech errors. 2009. xiii, 372 pp.
- 136 **KANDYBOWICZ, Jason:** The Grammar of Repetition. Nupe grammar at the syntax-phonology interface. 2008. xiii, 168 pp.
- 135 **LEWIS, William D., Simin KARIMI, Heidi HARLEY and Scott O. FARRAR (eds.):** Time and Again. Theoretical perspectives on formal linguistics. In honor of D. Terence Langendoen. 2009. xiv, 265 pp.
- 134 **ARMON-LOTEM, Sharon, Gabi DANON and Susan ROTHSTEIN (eds.):** Current Issues in Generative Hebrew Linguistics. 2008. vii, 393 pp.
- 133 **MACDONALD, Jonathan E.:** The Syntactic Nature of Inner Aspect. A minimalist perspective. 2008. xv, 241 pp.
- 132 **BIBERAUER, Theresa (ed.):** The Limits of Syntactic Variation. 2008. vii, 521 pp.
- 131 **DE CAT, Cécile and Katherine DEMUTH (eds.):** The Bantu-Romance Connection. A comparative investigation of verbal agreement, DPs, and information structure. 2008. xix, 355 pp.
- 130 **KALLULLI, Dalina and Liliane TASMOWSKI (eds.):** Clitic Doubling in the Balkan Languages. 2008. ix, 442 pp.
- 129 **STURGEON, Anne:** The Left Periphery. The interaction of syntax, pragmatics and prosody in Czech. 2008. xi, 143 pp.
- 128 **TALEGHANI, Azita H.:** Modality, Aspect and Negation in Persian. 2008. ix, 183 pp.
- 127 **DURRLEMAN-TAME, Stephanie:** The Syntax of Jamaican Creole. A cartographic perspective. 2008. xii, 190 pp.
- 126 **SCHÄFER, Florian:** The Syntax of (Anti-)Causatives. External arguments in change-of-state contexts. 2008. xi, 324 pp.
- 125 **ROTHSTEIN, Björn:** The Perfect Time Span. On the present perfect in German, Swedish and English. 2008. xi, 171 pp.
- 124 **IHSANE, Tabea:** The Layered DP. Form and meaning of French indefinites. 2008. ix, 260 pp.
- 123 **STOYANOVA, Marina:** Unique Focus. Languages without multiple wh-questions. 2008. xi, 184 pp.
- 122 **OOSTERHOF, Albert:** The Semantics of Generics in Dutch and Related Languages. 2008. xviii, 286 pp.
- 121 **TUNGSETH, Mai Ellin:** Verbal Prepositions and Argument Structure. Path, place and possession in Norwegian. 2008. ix, 187 pp.

- 120 ASBURY, Anna, Jakub DOTLAČIL, Berit GEHRKE and Rick NOUWEN (eds.): Syntax and Semantics of Spatial P. 2008. vi, 416 pp.
- 119 FORTUNY, Jordi: The Emergence of Order in Syntax. 2008. viii, 211 pp.
- 118 JÄGER, Agnes: History of German Negation. 2008. ix, 350 pp.
- 117 HAUGEN, Jason D.: Morphology at the Interfaces. Reduplication and Noun Incorporation in Uto-Aztecán. 2008. xv, 257 pp.
- 116 ENDO, Yoshio: Locality and Information Structure. A cartographic approach to Japanese. 2007. x, 235 pp.
- 115 PUTNAM, Michael T.: Scrambling and the Survive Principle. 2007. x, 216 pp.
- 114 LEE-SCHOENFELD, Vera: Beyond Coherence. The syntax of opacity in German. 2007. viii, 206 pp.
- 113 EYTHÓRSSON, Thórhallur (ed.): Grammatical Change and Linguistic Theory. The Rosendal papers. 2008. vi, 441 pp.
- 112 AXEL, Katrin: Studies on Old High German Syntax. Left sentence periphery, verb placement and verb-second. 2007. xii, 364 pp.
- 111 EGUREN, Luis and Olga FERNÁNDEZ-SORIANO (eds.): Coreference, Modality, and Focus. Studies on the syntax-semantics interface. 2007. xii, 239 pp.
- 110 ROTHSTEIN, Susan (ed.): Theoretical and Crosslinguistic Approaches to the Semantics of Aspect. 2008. viii, 453 pp.
- 109 CHOCANO, Gema: Narrow Syntax and Phonological Form. Scrambling in the Germanic languages. 2007. x, 333 pp.
- 108 REULAND, Eric, Tanmoy BHATTACHARYA and Giorgos SPATHAS (eds.): Argument Structure. 2007. xviii, 243 pp.
- 107 CORVER, Norbert and Jairo NUNES (eds.): The Copy Theory of Movement. 2007. vi, 388 pp.
- 106 DEHÉ, Nicole and Yordanka KAVALOVA (eds.): Parentheticals. 2007. xii, 314 pp.
- 105 HAUMANN, Dagmar: Adverb Licensing and Clause Structure in English. 2007. ix, 438 pp.
- 104 JEONG, Youngmi: Applicatives. Structure and interpretation from a minimalist perspective. 2007. vii, 144 pp.
- 103 WURFF, Wim van der (ed.): Imperative Clauses in Generative Grammar. Studies in honour of Frits Beukema. 2007. viii, 352 pp.
- 102 BAYER, Josef, Tanmoy BHATTACHARYA and M.T. Hany BABU (eds.): Linguistic Theory and South Asian Languages. Essays in honour of K. A. Jayaseelan. 2007. x, 282 pp.
- 101 KARIMI, Simin, Vida SAMIIAN and Wendy K. WILKINS (eds.): Phrasal and Clausal Architecture. Syntactic derivation and interpretation. In honor of Joseph E. Emonds. 2007. vi, 424 pp.
- 100 SCHWABE, Kerstin and Susanne WINKLER (eds.): On Information Structure, Meaning and Form. Generalizations across languages. 2007. vii, 570 pp.
- 99 MARTÍNEZ-GIL, Fernando and Sonia COLINA (eds.): Optimality-Theoretic Studies in Spanish Phonology. 2007. viii, 564 pp.
- 98 PIRES, Acrísio: The Minimalist Syntax of Defective Domains. Gerunds and infinitives. 2006. xiv, 188 pp.
- 97 HARTMANN, Jutta M. and László MOLNÁRFI (eds.): Comparative Studies in Germanic Syntax. From Afrikaans to Zurich German. 2006. vi, 332 pp.
- 96 LYNGBELT, Benjamin and Torgir SOLSTAD (eds.): Demoting the Agent. Passive, middle and other voice phenomena. 2006. x, 333 pp.
- 95 VOGELER, Svetlana and Liliane TASMOWSKI (eds.): Non-definiteness and Plurality. 2006. vi, 358 pp.
- 94 ARCHE, María J.: Individuals in Time. Tense, aspect and the individual/stage distinction. 2006. xiv, 281 pp.
- 93 PROGOVAC, Ljiljana, Kate PAESANI, Eugenia CASIELLES and Ellen BARTON (eds.): The Syntax of Nonsententials. Multidisciplinary perspectives. 2006. x, 372 pp.
- 92 BOECKX, Cedric (ed.): Agreement Systems. 2006. ix, 346 pp.
- 91 BOECKX, Cedric (ed.): Minimalist Essays. 2006. xvi, 399 pp.
- 90 DALMI, Gréte: The Role of Agreement in Non-Finite Predication. 2005. xvi, 222 pp.
- 89 VELDE, John R. te: Deriving Coordinate Symmetries. A phase-based approach integrating Select, Merge, Copy and Match. 2006. x, 385 pp.
- 88 MOHR, Sabine: Clausal Architecture and Subject Positions. Impersonal constructions in the Germanic languages. 2005. viii, 207 pp.
- 87 JULIEN, Marit: Nominal Phrases from a Scandinavian Perspective. 2005. xvi, 348 pp.
- 86 COSTA, João and Maria Cristina FIGUEIREDO SILVA (eds.): Studies on Agreement. 2006. vi, 285 pp.

- 85 MIKKELSEN, Line: Copular Clauses. Specification, predication and equation. 2005. viii, 210 pp.
- 84 PAFEL, Jürgen: Quantifier Scopepe in German. 2006. xvi, 312 pp.
- 83 SCHWIEKERT, Walter: The Order of Prepositional Phrases in the Structure of the Clause. 2005. xii, 338 pp.
- 82 QUINN, Heidi: The Distribution of Pronoun Case Forms in English. 2005. xii, 409 pp.
- 81 FUSS, Eric: The Rise of Agreement. A formal approach to the syntax and grammaticalization of verbal inflection. 2005. xii, 336 pp.
- 80 BURKHARDT SCHUMACHER, Petra: The Syntax–Discourse Interface. Representing and interpreting dependency. 2005. xii, 259 pp.
- 79 SCHMID, Tanja: Infinitival Syntax. Infinitivus Pro Participio as a repair strategy. 2005. xiv, 251 pp.
- 78 DIKKEN, Marcel den and Christina TORTORA (eds.): The Function of Function Words and Functional Categories. 2005. vii, 292 pp.
- 77 ÖZTÜRK, Balkız: Case, Referentiality and Phrase Structure. 2005. x, 268 pp.
- 76 STAVROU, Melita and Arhonto TERZI (eds.): Advances in Greek Generative Syntax. In honor of Dimitra Theophanopoulou-Kontou. 2005. viii, 366 pp.
- 75 DI SCIULLO, Anna Maria (ed.): UG and External Systems. Language, brain and computation. 2005. xviii, 398 pp.
- 74 HEGGIE, Lorie and Francisco ORDÓÑEZ (eds.): Clitic and Affix Combinations. Theoretical perspectives. 2005. viii, 390 pp.
- 73 CARNIE, Andrew, Heidi HARLEY and Sheila Ann DOOLEY (eds.): Verb First. On the syntax of verb-initial languages. 2005. xiv, 434 pp.
- 72 FUSS, Eric and Carola TRIPS (eds.): Diachronic Clues to Synchronic Grammar. 2004. viii, 228 pp.
- 71 GELDEREN, Elly van: Grammaticalization as Economy. 2004. xvi, 320 pp.
- 70 AUSTIN, Jennifer R., Stefan ENGELBERG and Gisa RAUH (eds.): Adverbials. The interplay between meaning, context, and syntactic structure. 2004. x, 346 pp.
- 69 KISS, Katalin É. and Henk van RIEMSDIJK (eds.): Verb Clusters. A study of Hungarian, German and Dutch. 2004. vi, 514 pp.
- 68 BREUL, Carsten: Focus Structure in Generative Grammar. An integrated syntactic, semantic and intonational approach. 2004. x, 432 pp.
- 67 MIŠESKA TOMIĆ, Olga (ed.): Balkan Syntax and Semantics. 2004. xvi, 499 pp.
- 66 GROHMANN, Kleanthes K.: Prolific Domains. On the Anti-Locality of movement dependencies. 2003. xvi, 372 pp.
- 65 MANNINEN, Satu Helena: Small Phrase Layers. A study of Finnish Manner Adverbials. 2003. xii, 275 pp.
- 64 BOECKX, Cedric and Kleanthes K. GROHMANN (eds.): Multiple *Wh*-Fronting. 2003. x, 292 pp.
- 63 BOECKX, Cedric: Islands and Chains. Resumption as stranding. 2003. xii, 224 pp.
- 62 CARNIE, Andrew, Heidi HARLEY and MaryAnn WILLIE (eds.): Formal Approaches to Function in Grammar. In honor of Eloise Jelinek. 2003. xii, 378 pp.
- 61 SCHWABE, Kerstin and Susanne WINKLER (eds.): The Interfaces. Deriving and interpreting omitted structures. 2003. vi, 403 pp.
- 60 TRIPS, Carola: From OV to VO in Early Middle English. 2002. xiv, 359 pp.
- 59 DEHÉ, Nicole: Particle Verbs in English. Syntax, information structure and intonation. 2002. xii, 305 pp.
- 58 DI SCIULLO, Anna Maria (ed.): Asymmetry in Grammar. Volume 2: Morphology, phonology, acquisition. 2003. vi, 309 pp.
- 57 DI SCIULLO, Anna Maria (ed.): Asymmetry in Grammar. Volume 1: Syntax and semantics. 2003. vi, 405 pp.
- 56 COENE, Martine and Yves D’HULST (eds.): From NP to DP. Volume 2: The expression of possession in noun phrases. 2003. x, 295 pp.
- 55 COENE, Martine and Yves D’HULST (eds.): From NP to DP. Volume 1: The syntax and semantics of noun phrases. 2003. vi, 362 pp.
- 54 BAPTISTA, Marlyse: The Syntax of Cape Verdean Creole. The Sotavento varieties. 2003. xxii, 294 pp. (incl. CD-rom).
- 53 ZWART, Jan-Wouter and Werner ABRAHAM (eds.): Studies in Comparative Germanic Syntax. Proceedings from the 15th Workshop on Comparative Germanic Syntax (Groningen, May 26–27, 2000). 2002. xiv, 407 pp.